

INTELLIGENT SYSTEMS QUALIFIER

Spring 2011

Each IS student has two specialty areas. Answer all 3 questions in each of your specialty areas.

You will be assigned an identifying number and are required to hand in printed or written copies of your answers with each page identified only by that number. This process enables us to grade your answers anonymously. You should NOT identify yourself explicitly by name on your answer sheets or implicitly by referring in the first person to your work (my project on ABC).

Please answer each question starting on a new page, with your answer following the text of the question.

Place any relevant references that you cite in an answer at the end of that answer, NOT in a separate section at the end of the whole exam.

If you have any questions or feel it necessary to make any assumptions in your answers, do not seek clarification from faculty or staff. Simply record your assumptions as part of your answer.

If one of your two areas is **Robotics**, answer the three questions below:

Robotics #1: There is a tension (both historically and intellectually) regarding the role of knowledge in behavior-based systems. Dan Dennett has characterized knowledge as taking three different representational forms: explicit, implicit, and tacit. There may be other useful taxonomies as well, but use the different qualities of knowledge to answer the questions below.

a) The claim that a truly reactive architecture is completely free of knowledge representation has served as the hallmark characteristic of nouvelle AI (as opposed to KBAI). Is it really true though? How according to Dennett's (or some other) taxonomy can knowledge be viewed as being represented even within a system such as the subsumption architecture? Or is it really not at all?

b) Payton, Borenstein and many others have added the use of knowledge representations that still permit the use of a behavior-based reactive system design, but still allow for explicit knowledge usage. How can this be? Why bother? What advantages does it afford? What consequences can occur if poorly managed?

c) Can you show and discuss how SLAM and behavior-based architectures can be conjoined, perhaps in the context of a hybrid deliberative/reactive architecture? Even if they can, should they be? Why or why not?

Robotics #2: Many robotics researchers report results from simulation in their papers and articles. Some reviewers claim that such results are not valid as "robotics" research. Summarize the main points for and against simulation in robotics research. Are there areas of research for which such results are valid, and others where they are not? Give concrete examples to back up your arguments and conclude with your views on the topic.

Robotics #3: In robotics frequently considered architectures include subsumption, behavior based, and more generally hybrid deliberative systems.

a) Characterize each of the three architectures in terms of underlying control, reliance on sensors and their arbitration mechanisms.

b) Describe scenarios in which each of the three architectures is particularly suited.

c) For research in modern service robotics using robots such as Simon, Jeeves, Golem-Krang, which of the architectures is best suited? Motivate the choice in terms of advantages and disadvantages.

If one of your two areas is **Perception**, answer the three questions below:

Perception #1: In structure from motion from unstructured internet collections, one of the most difficult problems is that of establishing correspondences between different images.

a) Many groups use a processing pipeline which includes (1) vocabulary trees, (2) RANSAC, (3) fundamental matrix. Please describe each of these three technologies and (i) why (ii) when (iii) how they are used within the SfM pipeline. Highlight how the product of each process is used.

b) An idea for a new paper is to use the trifocal tensor instead. What will change in the pipeline when we do that, and what are the issues that you will encounter? Will it work just as well, as fast, as robustly?

Perception #2: Formulate the perceptual grouping problem as a graph-partitioning problem. In your formulation also construct how this makes the problem a global segmentation problem as opposed to being driven by local regions. Then relate briefly how algorithms like Normalized Cuts, Min Cuts, and Graph Cuts are (collectively or individually) brought to bear on this problem. Do provide short descriptions of each, and their relationships to each other. Why do folks who like these graph-based approaches believe they are better than MRF based approaches? Why do folks who like these graph-based approaches believe they are better than other inference algorithms for MRF models, such as mean field approximation, loopy belief propagation, or MCMC?

Perception #3: The field of action recognition has introduced concepts like "Temporal-Templates" and "Space-Time Shapes." as representations. Describe these representations (and their formulations) and compare them. Is there something common between them? Further work in this area has taken these representations and then used other information to help with recognition of actions. What are the limitations of these representations and what additional information can help with action recognition in real videos (not just jumping jacks!)? What other spatio-temporal information needs to be added to these representations to really make these approaches work in real domains? Use these concepts to come up with a list of what measures should be included in any new forms of spatio-temporal features.

If one of your two areas is **Machine Learning**, answer the three questions below:

Machine Learning #1: The prisoner's dilemma is a classic and fundamental problem in game theory that demonstrates why two people might not cooperate even if it is in their best interests to do so:

Two suspects are arrested by the police. The police do not have enough evidence to convict them, so decide to convince them to admit guilt. They separate the two and offer each the same deal. If one testifies against the other (defects) while the other remains silent (cooperates), the defector goes free and the cooperator receives a ten-year sentence. If both cooperate, both are sentenced to only six months. If each defects, each receives a five-year sentence.

The basic question is: as one of the prisoners, how should one act? It can be shown that the Nash equilibrium is for both prisoners to defect, leading to what we might think of as a fairly undesirable outcome, at least for the prisoners.

Now imagine that the prisoners will face this situation with each other multiple times. Further, imagine that they have one step of memory.

Construct a multi-agent decision problem and show how the agents will learn how to act over time. In particular, argue for what they will learn to do. Does your answer violate Nash? Why or why not?

Machine Learning #2: In designing Markov Decision Processes (MDPs), one must define the world (as states), its dynamics (as a transition model), a set of actions, and a reward function. The solution to an MDP is a policy that maps states to actions. The optimal policy maximizes long term expected reward.

Imagine a different formulation where we are interested in building adaptive games. In this problem, trajectories of states matter, not just states. Further, we want to allow for replayability; that is, a player who plays the game multiple times should experience a variety of play even if she takes the same actions. In this formulation, we will redefine reward to be a distribution over possible trajectories

through the space; therefore, a policy is now non-deterministic, mapping partial trajectories to distributions over actions.

Define the analogous changes (if any) to the other components of an MDP. In particular, define precisely what an optimal policy is in this case. Sketch how you would construct an algorithm to find such an optimal policy. Is your algorithm computationally feasible? If not, what approximations will you make?

Machine Learning #3: Machine learning algorithms have traditionally had difficulty scaling to large problems. In classification and traditional supervised learning this problem arises with data that exist in very high dimensional spaces or when there are many data points for computing, for example, estimates of conditional densities. In reinforcement learning this is also the case, arising when, for example, there are many, many states or when actions are at a very low level of abstraction.

a) Typical approaches to addressing such problems in RL include function approximation and problem decomposition. Compare and contrast these two approaches. What problems of scale do these approaches address? What are their strengths and weaknesses? Are they orthogonal approaches? Can they work well together?

b) What are the differences between hierarchical and modular reinforcement learning? Explain both the theoretical and practical limits of these approaches.

If one of your two areas is **Planning and Search**, answer the three questions below:

Planning and Search #1: Heuristic search is commonly accepted as a leading method for solving planning problems in both classical and motion planning domains.

a) Define what it means for a heuristic to be admissible. Why is it desirable for a heuristic to be admissible? Is it always desirable? If so, explain why. If not, present an example domain and non-admissible heuristic and explain the benefit.

b) What is the purpose of the Fast Forward (FF) heuristic? What kind of domains does it apply to? Is this heuristic admissible? Explain your answer.

c) Many modern planners, including FF use best-first search rather than A*. Explain the expected effect this has in terms of the completeness, efficiency and optimality of the planner. Point out any specific problems that may occur during planning. Discuss possible strategies that a best-first planner could use to overcome these problems.

Planning and Search #2: Consider creating a motion planner for an autonomous car-like vehicle. The vehicle must quickly travel to a target location while avoiding various obstacles such as buildings and rubble. A map of the terrain is given to you ahead of time.

a) Rapidly Exploring Random Tree (RRT) is a popular algorithm in motion planning. Is this a good choice for this domain? Discuss the pros and cons of applying this algorithm to the problem at hand. Include in your discussion the specific variant of RRTs you might consider using as well as arguments with regard to the completeness, optimality and efficiency of the resulting plan.

b) Instead of using RRTs, suppose you approach the problem using cell decomposition. Choose two cell decomposition algorithms (one exact and one approximate). Compare and contrast these algorithms for this domain. Explain which algorithm you would prefer to use and justify your reasoning.

c) What is the primary challenge for cell decomposition methods in this domain? Is there a way to resolve this challenge by incorporating RRTs or other sampling based planners? Discuss how you might do this or explain why not. Similarly, is there a way to resolve it without incorporating any randomized strategy? Explain your answer.

Planning and Search #3: Classical planners and motion/path planners are useful tools that apply to different types of domains. However, many real-world automation challenges do not fit squarely into either category.

a) Present an example domain that could be solved by a combination of classical and path planning but is difficult to represent in terms of one of these formulations alone. Discuss why this is the case for each type of planning.

b) Come up with a planning algorithm that combines two planners, one from each category (classical and motion) to solve your problem. Define the domain for each planner. Be specific about which planning methods you chose, explain why you chose them and how they are combined.

c) What are the challenges and anticipated problems in combining your planners? How do you propose to address these challenges?

If one of your two areas is **Cognitive Science**, answer the three questions below:

Cognitive Science #1: On the traditional cognitive science account, culture is content for the mind to operate on; that is, what is important about culture for cognition is abstracted in the representations used in various cognitive processes. Contemporary critics argue that cognition itself is inherently cultural; that is, culture and cognition are co-constitutive. Hutchins (1995), Shore (1996) , Tomasello (1999) all argue for view but from different perspectives, using different kinds of evidence.

- a) Briefly develop each account.
- b) In what ways are they compatible/incompatible?
- c) Should a committed traditionalist be swayed?

Cognitive Science #2: Historically, cognitive science has sometimes considered action and perception as secondary to cognition. But this view has changed dramatically over the last generation, and now most cognitive scientists recognize that action and perception are fundamental and critical parts of cognition. This new view shows up in many theories of cognition ranging from enactive representations to visual thinking to embodied cognition. Take any one example of enactive representation or visual thinking or embodied cognition in cognitive science, and describe how the new view differs from the old one. What are the advantages of the new view?

Cognitive Science #3: Recently, Bill Murdock gave a talk about the IBM Watson system that defeated human champions in the game of Jeopardy. This victory follows IBM's Deep Blue victory over Gary Kasparov in the game of chess. *Bring your readings to bear on your answers to the following:*

- a) What did IBM Deep Blue tell us about human cognition?
- b) What does IBM Watson tell us about human cognition? In what ways are they different?

If one of your two areas is **Knowledge Based AI**, answer the three questions below:

Knowledge Based AI #1: According to some AI researchers, a fundamental conundrum of AI is this: Most interesting AI problems have global constraints, but computation is fundamentally local. Clearly, no one wants to build an AI that can understand only a part of a story; we want an AI that can understand the complete story. Similarly, no one wants to build an AI that can design only a part of a system; we want an AI that can design the complete system. Yet, any AI must understand a story or design a system, one small part at a time (even if one makes the process concurrent).

Describe three principles from knowledge-based AI that seek to address this conundrum. Illustrate these principles with examples (including examples from your own work).

Knowledge Based AI #2: Recently Ben Shneiderman gave a GVU talk on information visualization techniques and tools for knowledge discovery. It is noteworthy that his HCI laboratory developed many of the techniques and tools he described initially in the context of supporting human creativity. Instead of HCI, let us take the perspective of KBAI to the task of supporting human creativity. Describe and illustrate three design principles or methods for supporting human creativity from the AI perspective. Argue why and how your principles are right/useful.

Knowledge Based AI #3: Discuss the relationship between each of the following pairs of reasoning methods. For each pair, explain whether or not these are different names for the same method. Are they subsets of one another? Are they entirely different concepts with clear boundaries? Include in your discussion the similarities and differences between the methods from both the perspective of theory and implementation as tools for computational reasoning. Provide example problems that are best suited to a particular method and explain why these problems clearly distinguish one from another.

- a) Analogical Reasoning and Conceptual Blending
- b) Analogical Reasoning and Case-Based Reasoning

You can learn more about conceptual blending at
<http://www.cogsci.ucsd.edu/~faucon/BEIJING/blending.pdf>