INTELLIGENT SYSTEMS QUALIFIER

Spring 2016

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 areas is **Perception**, answer the three questions below:

Perception #1

UAV: Recently Unmanned Aerial Vehicles (UAVs) have emerged as very promising platforms for a variety of robotics applications, including powerline inspection, surveillance, and so forth. UAVs which can carry a color video camera create new opportunities for vision-based surveillance and control. Your job is to design a vision-based surveillance system for a UAV which will enable it to follow a single person and recognize what they are doing, as part of an industrial security system. The question has two parts:

a) Visual Tracking
Briefly outline the challenges that are involved in tracking a moving person using a moving camera. What assumptions could you make to constrain the problem and how would they help? What sources of information in addition to the video could you make use of, and how would they help? Briefly describe an algorithm for person tracking in this situation and identify the conditions under which it would work and some of its failure modes (when it would fail to track reliably).

b) Visual SLAM
Now imagine that you also have a visual SLAM system running on the UAV. How could the visual SLAM system be used to improve the performance of the tracker? Briefly describe one possible SLAM algorithm for this situation. Describe the components of your SLAM method and draw a block diagram to show how they fit together. What are the most challenging parts of the SLAM problem? How is SLAM on a UAV different from standard ground vehicle applications of SLAM?

Perception #2

Fakeblock: You are hired by Twitter to stop the flood of dishonest photoshopped pictures being disseminated, e.g. https://en.wikipedia.org/wiki/Kerry_Fonda_2004_election_photo_controversy . In particular, you are looking for photos which are the composite of pieces of two real photos. There are many existing approaches to this task based on low level signal processing (e.g. JPEG statistics being unusual) or lighting (e.g. shadows falling in inconsistent directions) but you are tasked with taking a recognition approach.

a) You will need to train a classifier from training examples which can detect composite images. Describe how you will assemble a sufficient dataset for training and testing. What properties of a composite image would you expect your classifier to be able to exploit? Describe the features and classifiers you would use in your learning approach.

b) Give some examples of properties of composite images that would make them easy to detect. Give an example of a composite that would be difficult to detect. (Note that these two questions *are not* asking you to provide image examples, we are asking you to describe the properties of easy and hard cases in words.) Can you identify applications of your detector in other disciplines, such as computer graphics?

Perception #3

You are part of a team that is collecting videos from different cities while walking around or driving on the street. The goal is to find signs (Billboards, posters, and other visuals that are of a political nature). However, due to some reason, it is required that in the video shown to public, all street numbers, street names, addresses, and car licence plates numbers are blurred out, exception to that is the name/model number of the car and other signs that are not "privacy sensitive". Your task is to implement an algorithm to blur numbers and text in this video, but as noted above, not all the text and numbers.

1. Suggest methods you think you'd like to use for this task. To keep it focused, we suggest you identify TWO methods/papers (they could be from the reading list, but don't have to be). It is important that these TWO are really different approaches (really different, one is detection based, other is tracking based, so both should be different).

2. Justify your choice, and highlight the advantages and disadvantages of each.

3. Highlight their differences.

4. Be specific about how your method can be used to blur only street names, numbers, and license plates but not other text in the video (signs, posters, car names, etc.).

5. Would Deep Learning/CNNs be a good method for this. Why? In brief, set-up the problem for CNNs.

6. How would one evaluate such a system?

-------------------------------------------------------------------------------------------------------------------

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

KBAI #1

Terms like "cognitive systems" and "cognitive computing" are in vogue again. IBM, for example, claims that Watson is a cognitive system that has ushered in a new era of cognitive computing. The design of Watson is still evolving. For the purposes of this question, here is how one version of Watson available to academia, called the Watson Engagement Manager (WEM), works. (1) A knowledge engineer seeds WEM with a corpus of natural language documents. The engineer segments each document into sections and headings; WEM understands the difference between them. These documents are WEM's memory of answers. (2) The knowledge engineer trains WEM by building a taxonomy of questions, and identifying specific section headings in specific documents that best answer each type of question. These question types are WEM's indices into its answer memory. (3) The knowledge engineer further trains WEM by building variations of each type of question and clustering the variations around the question type. (4) When a user poses a question to WEM in natural language, WEM identifies the keywords in the question through lexical processing. (5) WEM finds the question type that most closely matches the vector of keywords in the question through vector matching. (6) WEM retrieves all candidate answers (section headings and the following text) from its memory. (7) WEM calculates confidence scores for each candidate answer based on the frequency of occurrence of the keywords in the answer (actual text). (8) Watson selects the answer with the highest confidence value; if the confidence value is higher than some preset threshold, it returns the answer to the user. (Again, for the purposes of the question, assume that this is how WEM works; you do not need to chase the literature on Watson.)

(a) Mainstream AI scientists reject IBM's claims: "WEM does not work like the human mind," they say. In what ways, if any, does WEM work like the human mind? What abstractions in WEM, if any, have cognitive correlates?

(b) IBM counters that while Watson may or may not work like the human mind, it helps humans address complex problems that require higher-level cognition, and it is in this sense that Watson is a cognitive system. How might one decide whether or not even this hypothesis is valid? Give a high-level design of an experiment that may help answer this question.

KBAI #2

This question refers to the following story by Isaac Asimov. Let us assume (just for the purposes of this examination) that the knowledge-based theories of intelligence are accurate and precise portrayals of human cognition. Thus, Willm Shakesper (let us suppose) actually does have production rules and semantic networks and frames in his mind, uses means ends analysis and problem reduction to address problems, relies on case-based reasoning and incremental concept learning to interact with world, etc. Yet, according to Asimov, Willm Shakesper, and Archimedes and Newton and Galileo, all fail in our modern world. Willm Shakesper even flunks the course on William Shakespeare. This means while these very intelligent people may have all the knowledge-based methods, somehow their methods were not very useful to deal with the modern world. The question is why?

a: Let us suppose that Willm Shakesper has many semantic networks in his mind. Why might

his semantic networks not work in the modern world?

b: Let us suppose that Willm Shakesper knows about and can execute means ends analysis. Why might his means ends analysis method not work in the modern world?

c: Let us suppose that Willm Shakesper has many production rules in his mind. Why might his production system not work in the modern world?

d. Let us suppose that Willm Shakesper has many cases in his mind. Why might his case application not work well in the modern world?

e: Let us suppose Willm Shakesper can learn new concepts via incremental concept learning. Why might even incremental concept learning not work very well in the modern world?

## THE IMMORTAL BARD

*by Isaac Asimov*

"Oh, yes," said Dr. Phineas Welch, "I can bring back the spirits of the illustrious dead."

He was a little drunk, or maybe he wouldn't have said it. Of course, it was perfectly all right to get a little drunk at the annual Christmas party.

Scott Robertson, the school's young English instructor, adjusted his glasses and looked to right and left to see if they were overheard. "Really, Dr. Welch."

"I mean it. And not just the spirits. I bring back the bodies, too."

"I wouldn't have said it were possible," said Robertson primly.

"Why not? A simple matter of temporal transference."

"You mean time travel? But that's quite-uh-unusual."

"Not if you know how."

"Well, how, Dr. Welch?"

"Think I'm going to tell you?" asked the physicist gravely. He looked vaguely about for another drink and didn't find any. He said, "I brought quite a few back. Archimedes, Newton, Galileo. Poor fellows."

"Didn't they like it here? I should think they'd have been fascinated by our modern science," said Robertson. He was beginning to enjoy the conversation.

"Oh, they were. They were. Especially Archimedes. I thought he'd go mad with joy at first after I explained a little of it in some Greek I'd boned up on, but no-no-"

"What was wrong?"

"Just a different culture. They couldn't get used to our way of life. They got terribly lonely and frightened. I had to send them back."

"That's too bad."

"Yes. Great minds, but not flexible minds. Not universal. So I tried Shakespeare."

"*What*?" yelled Robertson. This was getting closer to home.

"Don't yell, my boy," said Welch. "It's bad manners."

"Did you say you brought back Shakespeare?"

"I did. I needed someone with a universal mind; someone who knew people well enough to be able to live with them centuries way from his own time. Shakespeare was the man. I've got his signature. As a memento, you know."

"On you?" asked Robertson, eyes bugging.

"Right here." Welch fumbled in one vest pocket after another. "Ah, here it is."

A little piece of pasteboard was passed to the instructor. On one side it said: "L. Klein & Sons, Wholesale Hardware." On the other side, in straggly script, was written, "Will$_m$ Shakesper."

A wild surmise filled Robertson. "What did he look like?"

"Not like his pictures. Bald and an ugly mustache. He spoke in a thick brogue. Of course, I did my best to please him with our times. I told him we thought highly of his plays and still put them on the boards. In fact, I said we thought they were the greatest pieces of literature in the English language, maybe in any language."

"Good. Good," said Robertson breathlessly.

"I said people had written volumes of commentaries on his plays. Naturally he wanted to see one and I got one for him from the library."

"And?"

"Oh, he was fascinated. Of course, he had trouble with the current idioms and references to events since 1600, but I helped out. Poor fellow. I don't think he ever expected such treatment. He kept saying, 'God ha' mercy! What cannot be racked from words in five centuries? One could wring, methinks, a flood from a damp clout!'"

"He wouldn't say that."

"Why not? He wrote his plays as quickly as he could. He said he had to on account of the deadlines. He wrote *Hamlet* in less than six months. The plot was an old one. He just polished it up."

"That's all they do to a telescope mirror. Just polish it up," said the English instructor indignantly.

The physicist disregarded him. He made out an untouched cocktail on the bar some feet away and sidled toward it. "I told the immortal bard that we even gave college courses in Shakespeare."

"*I* give one."

"I know. I enrolled him in your evening extension course. I never saw a man so eager to find out what posterity thought of him as poor Bill was. He worked hard at it."

"You enrolled William Shakespeare in my course?" mumbled Robertson. Even as an alcoholic fantasy, the thought staggered him. And *was* it an alcoholic fantasy? He was beginning to recall a bald man with a queer way of talking....

"Not under his real name, of course," said Dr. Welch. "Never mind what he went under. It was a mistake, that's all. A big mistake. Poor fellow." He had the cocktail now and shook his head at it.

"Why was it a mistake? What happened?"

"I had to send him back to 1600," roared Welch indignantly. "How much humiliation do you think a man can stand?"

"What humiliation are you talking about?"

Dr. Welch tossed off the cocktail. "Why, you poor simpleton, you *flunked* him."


KBAI #3

This question refers to the following story by Brian Baer. In this story, the protagonist performs a kind of diagnosis and repair of Henry the robot.


(a) Illustrate the production system architecture for the diagnosis task in the story. Invent percepts, actions and rules as needed (with at least 3 percepts, 3 actions, 3 rules). Demonstrate how the production system architecture would work for this story (that is, given Henry's symptoms, it would ask the same questions that our protagonist asks). Show the evolving contents of the working memory.

(b) Illustrate the process of case-based reasoning for the diagnosis task in the story. Invent percepts, actions and cases as needed (with at least 3 percepts, 3 actions, 3 cases). Demonstrate how the case-based reasoning process would work for this story (that is, given Henry's symptoms, it would result in the same diagnosis as in the story). Show at least the steps of case retrieval and adaptation.

(c) Now suppose that one of the rules was missing in the long-term memory of the production system. Show why, when and how this rule may be learned from the case base.

**The Robot Whisperer**

**Brian C. Baer (2008)**


Robots love me.

As much as robots can love. And in a Platonic sense, of course. Something about my chubby little baby face sets off their simulated paternal instincts and they all bend over backwards to answer my questions. That sort of thing comes in handy with my job.

I knelt in front of the unmoving blue robot. As if brooding, it sat on the floor in the middle of the living room. It was large and bulky, a few years old but in decent enough shape. Not one of those smooth, humanoid-looking models that have been flooding the market. It was more from the "Rock 'Em, Sock 'Em" school of design. Behind me, the family stood anxious, worried, huddled together.

"Can you fix him, doctor?" the wife asked. She hugged her young daughter close, and her husband did the same to her.

"I'm not a doctor," I said absent-mindedly as I eyed my scanner.

"I beg your pardon?" the husband chimed in, brushing a loose strand of hair across his comb-over with his palm.

"Hm?" I asked, coming out of my focus. "Oh. I'm not a doctor. Robots don't really have brains, so they don't need a psychiatrist or anything like…" I trailed off, before looking back to my

work. "I'm a technician."

"Henry just sat down and stopped moving," said the little girl, sounding close to tears.

"We just had him in for maintenance and everything checked out," the wife added. "I don't understand it."

I nodded and made a little "hmm" sound, but I wasn't really listening. "Unit NX-6401, respond to my voice."

"Henry," the robot corrected me in a surprisingly human voice. It still hadn't moved, and the lights hadn't returned to its dim photoreceptors.

"Okay, Henry," I conceded. "Are you functioning correctly?"

It made a soft snorting noise. "If that's what you call this."

I sat cross-legged on the carpet in front of it. "Hey, now. What's that all about?" I put my hand on its shoulder. Henry's ocular lights activated, but just barely. It didn't respond right away.

"The Johnsons across the street bought a new robot," it said finally.

"Yeah," the husband confirmed from behind me, "One of those new A-01 models."

"Go on," I coaxed.

"I've seen it walking their kids to school and fixing their roof, and it's got those extendable arms and a hedge-clipper accessory, and..."

"And its making you feel not as special?" I asked in a soothing voice.

"The A-01s are so great," it said. "One of them would be so much more functional for this family. It would be better than I am."

"Henry, I'm going to tell you a secret about humans. It is a bit paradoxical, so promise me your head will not explode when I tell you."

It nodded, its eyes glowing brighter.

"Henry," I said. "Humans build emotional attachments. And they don't always want what's shiny and new. They want what they love."

"They *love* me?" it asked. It stood up, and after a moment, I followed.

"It isn't very logical, doctor." Henry's voice sounded happy.

I smiled. "I'm not a doctor."

(This story has been slightly edited.)

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

Machine Learning #1

Gibbs Sampling is a widely-applicable technique for Bayesian inference in graphical models.

(a) Briefly summarize Gibbs Sampling, explaining when it can be applied, what is needed to apply it, and what guarantees (if any) are available.

(b) When would you apply an alternative such as slice sampling or Metropolis-Hastings instead of Gibbs sampling?

(c) Consider the problem of inducing a discrete hidden Markov model, given data generated from a model with very strong probabilistic dependencies between adjacent hidden states. Why might Gibbs sampling perform poorly in this model? What alternative sampling algorithm might do better?

(d) Please answer the following two questions on Markov Random Fields and Gibbs Sampling:

West et al (2014, "Exploiting Social Network Structure for Person-to-Person Sentiment Analysis") describe a *signed social network* in which each edge $(i, j) \in E$ has a sign expressing the (latent) sentiment between the two nodes, $Y_{i,j} \in \{+, -\}$, which is revealed by some linguistic features of the edge, $X_{i,j}$. Assume the graph is symmetric, so that $y_{i,j} = y_{j,i}$, and $x_{i,j} = x_{j,i}$. The graph is not assumed to be connected.

Let's model this as a Markov Random Field. We have potential functions $\psi_{x|y}(x_{i,j}, y_{i,j})$ for all edges $(i, j) \in E$. West et al also compute a score for edge triads, which we can model as a potential $\psi_y(y_{i,j}, y_{j,k}, y_{i,k})$ for each triad $(i, j, k) \in \text{Triads}(E)$. We then have the conditional likelihood,

$$P(y \mid x; E) \propto \left( \prod_{(i,j) \in E} \psi_{x|y}(x_{i,j}, y_{i,j}) \right) \times \left( \prod_{(i,j,k) \in \text{Triads}(E)} \psi_y(y_{i,j}, y_{i,k}, y_{k,j}) \right). \tag{1}$$

You may assume that the triad potentials are invariant under rotation, e.g., $\psi_y(+, +, -) = \psi_y(+, -, +) = \psi_y(-, +, +)$.

(a) Given the social network $E = ((a, b), (b, c), (a, c), (a, d), (b, d), (b, e))$, draw the corresponding factor graph, including all random variables $x_{i,j}$ and $y_{i,j}$.

(b) West et al prove that exact inference in this model is intractable. They suggest Gibbs sampling as a possible approximation, although they do not explore this possibility in the paper. Derive a Gibbs sampling equation for individual edge signs. Your sampling equation should be expressed in terms of the $\psi(\cdot)$ functions defined above.

Machine Learning #2

Overfitting, underfitting, cross-validation, and multiclass problems:

(a) Please give pictorial examples in regression for overfitting and underfitting respectively, and explain the relation between over/underfitting and model class flexibility.

(b) We would like to perform k-fold cross-validation to select models. What should k be? Discuss the pros and cons of large or small values of k (in terms of bias, variance and computation).

(c) For ridge regression, how to perform leave-one-out cross-validation efficiently?

(d) Multiclass classification tries to assign one of several class labels (rather than binary labels) to an object. Give two ways which use binary classifiers to solve multi-class classification problem?

(e) What are the pros and cons of these different methods (e.g. in terms of computional complexity or the applicability of the method) of doing multiclass classification?

(f) Besides using binary classifiers, do you have any other idea on how to build a multiclass classifier?

(g) In activity recognition, often there is a NULL class that absorbs everything that is not of interest. For example, using an accelerometer in a smartwatch, we might want to classify fitness related events such as walking, running, biking, and eating and everything else is placed into a NULL class for training. Often times, this approach leads to many false positives. How might you address this problem, especially with respect to a-f above?

Machine Learning #3

Scaling up reinforcement learning:

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.

• 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?

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

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

Robotics #1

You are hired by a robotics company that manufactures supply delivery robots that are deployed in hospitals, hotels and office buildings. The robots act a mobile transport for goods within a given building, such as delivering items between the hotel front desk and individual hotel rooms. The navigation environment of the robots is challenging, with unexpected obstacles, reflective surfaces and other challenges found in the environment. As a result, the robots occasionally fail to navigate to their destination due to path planning errors. When this happens, the robots contact the company and one of the employees remotely teleoperates the robot until it is able to act autonomously again.

You have been hired to develop a navigation system that improves over time based on the teleoperation input provided each time an error occurs.

1. Describe the path planning technique you will use for the underlying navigations system. You can assume a map and localization system are provided.

2. Propose an algorithm that uses teleoperation input provided by a human operator to incrementally improve navigation performance within a given domain. For example, if the robot gets stuck/lost in a particular hallway multiple times and each time a human operator teleoperates it to correct the situation, how can your path planning representation improve over time in order to enable the robot to operate in the hallway autonomously in the future? Characterize the types of errors your algorithm can handle and how adaptation would be performed.

Robotics #2

"The world is its own best model" --R. Brooks

a) Explain what Brooks means by this statement and give an example robotics/control problem for which this is approach is used.

b) Explain the alternate viewpoint/approach and give an example robotics/control problem for which this approach is used.

Robotics #3

Over the last few years we have seen tremendous progress on robot motion planning from basic search to probabilistic methods. Given a configuration space for a particular problem, characterize when it is appropriate to use regular graph search, probabilistic roadmaps and rapidly exploring random trees. When are each of the methods most appropriate? What is the respective complexities and what are the key design trade-offs in their use across sensing, complexity of environment, …?

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

Cognitive Science #1

Appraisal theory is a cognitive account of emotion. In its simplest form, appraisal theory states that a negative emotion arises when some change in the world interferes with the attainment of a goal. Likewise a positive emotion arises when some change in the world makes goal attainment easier. For example, an entity might experience a negative emotion if it believes that a goal state has become harder to achieve or farther away, and an entity might experience a positive emotion if it believes that a goal state has become easier to achieve or closer than would otherwise be expected. The intensity of the emotional response is a function of the degree to which goal attainment is perceived to be harder or easier. Explain how Barsalou's Perceptual Symbol theory could be used to account for appraisal theory. Discuss how Perceptual Symbols might account for both negative and positive emotions. Your answer should account for the intensity of emotion and not just the presence or absence of an emotional response.

Cognitive Science #2

Holyoak and Thagard describe the ACME theory of analogical retrieval and mapping that used three kinds of constraints: structural, semantic, and pragmatic. In that same decade, Anderson developed the ACT-R architecture as a unified theory of cognition that represented basic symbolic knowledge representations and cognitive processes. Discuss the overlap between these two approaches. In particular, consider how ACME's constraints could be (or could not be) used in ACT-R.

Cognitive Science #3

You will graduate with a PhD in CS with specialization in Intelligent Systems and Cognitive Science. Both IS and Cogsci  consider theory as an important vehicle to move the field forward.

1. Consider two theories: Structured Mapping (Gentner 1983) and Situated cognition (Hutchins 1995). Compare and contrast the roles these two theories play in Cognitive Science.
2. What IS/CogSci theory most closely impacts your work?
3. Briefly present one of your recent projects.
   a. Consider how each of these theories (Structured Mapping, Situated Cognition, your favorite theory) can be used to
        i. generate a hypothesis
        ii. evaluate it in the context of your project.
   b. Is one of these theories more suited to the design of a study than the other?

---------------------------------------------------------------------------------------------------------------

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

Planning and Search #1

Design a grid world in which A* and D*Lite perform identically when re-planning in terms of number of states expanded. Your solution should show the grid world and indicate the following:

a) the perceptual threshold (e.g., can the agent see 1 cell away, 2 cells away, etc.)

b) the initial state, goal state, and the state in which re-planning occurs

c) which nodes are expanded by the initial A*, which nodes are expanded by the A* when re-planning occurs, and which nodes are expanded when D*Lite re-plans.

Alternatively, you may argue (or prove) that no grid world can exist where D*Lite replanning and A* perform identically.


Planning and Search #2

In Heuristic Search Planning (HSP), the domain-independent heuristic is computed by

$$h_a(p; s) = \begin{cases} 0 & \text{if } p \in s \\ \min_{a \in O(p)}[h_a(a; s)] & \text{otherwise} \end{cases}$$

where

$$h_a(a; s) = 1 + \sum_{q \in Pre(a)} h_a(q; s)$$

computed on a relaxed domain in which action delete lists are removed.

The Fast Forward (FF) planner uses Graphplan style search on a relaxed domain in which action delete lists are removed and generally computes the heuristic value of a state faster than HSP.

In both HSP and FF the cost of an action, *cost(a)* = 1 for all *a* in the set of operators, *O*.

a) Is the HSP heuristic with non-uniform cost(a) function admissible? Give a sketch of the proof or the intuition behind a proof.

b) Is the FF heuristic with non-uniform cost(a) function admissible? Give a sketch of the proof or the intuition behind a proof.

Planning and Search #3

In the last week, Google/DeepMind's AlphaGo system won three games of Go against the 3rd best human player, Lee SeDol. At the time this question is written, Lee SeDol in turn has won one game and there is one game yet to be played.

Go is a classic board game played by laying black or white stones on a 19 x 19 board with the goal of "capturing" opponent stones by completely encircling them with one's own stones. Go has a branching factor of approximately 250. Games can often last 100 moves or more. Contrast this with Chess, in which the branching factor is 35 and can last for a few dozen moves. Players of Go use a lot of metaphors to describe strategic plays, such as creating "living shapes" and "running away". Go experts extensively study the playing styles of great players throughout history.

AlphaGo combines two algorithmic techniques: Monte Carlo Tree Search (MCTS) and Deep Neural Networks. MCTS is conceptually similar to MiniMax, except instead of performing exhaustive search down to a depth limit, it randomly selects moves for both sides. The more sampling performed, the more accurate the estimate of a state's value. AlphaGo uses two deep neural nets. The first learns to estimate the value of board configurations. The second learns to predict which parts of the search space should be sampled. AlphaGo was trained for many months on historical traces of games. It also trains as it plays many hundreds of thousands of games against itself.

Commentators—expert Go players themselves—were taken by surprise several times in game #2 by moves made by AlphaGo, declaring them counterintuitive according to conventional strategies. AlphaGo won game #2. During game #4, Lee SeDol made a move that surprised commentators. AlphaGo never recovered and Lee SeDol won game #4.

a) Provide an argument for how AlphaGo is able to devise moves that take human expert players by surprise and yet are highly effective. Reference the technical aspects of AlphaGo that would be responsible (you don't need to research AlphaGo beyond what is described in this question).

b) Provide an argument for why Lee SeDol's surprise move in game #4 led to AlphaGo's downfall in that game. Reference the technical aspects of AlphaGo that would be responsible (you don't need to research AlphaGo beyond what is described in this question).