CS365: Artificial Intelligence

End-semester exam — Total Marks = 136, Time = 3 hours

Part A. **(Logic)**

Question 1. Consider the following : [4x5=20]

i. No student passing CS365 will fail to translate this sentence into logic.

ii. Alan cannot translate this sentence.

iii. Alan is a student.
Therefore, Alan has not passed CS365.
Notation: S(x): student, P(x): pass 365, T(x): translate sentence

(a) Tranlate the sentences into First-Order Logic

(b) Construct the proof using FOL.

(c) Translate the sentences to clause form

(d) Construct a proof using resolution refutation. Indicate unifying substitutions.

Question 2. Consider these two expressions in first-order logic: [4x5 = 20]

$$\text{(A)}: \ \forall x \, \exists y \ (x \geq y)$$
$$\text{(B)}: \ \exists y \, \forall x \ (x \geq y)$$

(a) Convert these sentences into clause form, using the predicate `geq(x,y)` for $x \geq y$.

(b) Use resolution to prove that if B holds then so must A.

(c) Standard Logic has four connectives - $\vee, \wedge, \Rightarrow, \sim$. In resolution, we have only $\vee, \sim$. Can you think of a single connective so that all the others can be defined in terms of it?

(d) "All men are mortal". Discuss how a baby eventually gets to learn general rules like this.

Part B. **(NLP)**

Question 3. Use the grammar E0 below for the next few parts [4x5=20]

```
grammar E0:
S  -> NP VP                    VP -> Verb
NP -> Pronoun                   | VP NP
| Noun                  | V S
| Article Noun          | VP Adverb
| NP PP                 | VP PP
| Adjective NP         PP -> Preposition
```

(a) Give an example of a) a sentence that is overgenerated, and b) a correct sentence that is rejected by this grammar.

(b) Generate at least two parses for the sentence: `Police help dog bite victim`

(c) Define a lexicon that can accommodate both your parses.

(d) In trying to disambiguate the above, can you assign probabilities to some of the rules of E0 so that one parse is more likely than the other? Explain what would it mean to disambiguate the sentence at i) the semantics level and ii) the pragmatic level.

Question 4. [4x5=20]

(a) By drawing a picture for it comment on whether one of the following sentences more difficult to draw than the other:
1. Colorless green ideas sleep furiously.
2. Furiously sleep ideas green colorless.
Does this imply that these sentences have a gradation in terms of semantic "validity"?

(b) Give the argument for the autonomy of syntax and suggest how (a) may impact it.

(c) Consider two different models of emotion recognition. One distinguishes between composed, agreeable, elated, confident, tired, confused, and the other between sad, fear, joy, guilt, shame, disgust, anger. How could **Latent Semantic Analysis** be used to map between these different vocabularies?

(d) What does it mean for a symbol to be "grounded"?

Part C. **(Vision / Robotics / Projects)**

Question 5. **Vision** [12x2=24]

i. Image forensics is based on using models of the image that consider a) high frequency components b) low frequency components c) ignore frequency

ii. In image tracking, you are given two sub-images as colour histograms. Define the "Bhattacharya distance" between these sub-images.

iii. What is a chaincode histogram?

iv. Explain how you may detect a "spatio-temporal interest point"

v. A Convolution neural network operates over a spatial neighbourhood. How can it be used for learning temporal sequences?

vi. How does kinect obtain depth data for every pixel?

vii. A common technique for object recognition and semantic scene analysis is "bag of visual words". Give an example of a visual "word" and how it may be computed.

viii. What is the difference between bottom-up and top-down attention?

ix. How is a model of the background learned for foreground extraction from static camera images.

x. How is the imaging for the Nao used to estimate the distance to a ball it has to grasp?

xi. Explain the similarities and dissimilarities between [2x2]
a) Adaboost as used in the Viola-Jones algorithm.
b) Random forest classifiers.

Question 5. **Robotics / Search / NLP** [16x2=32]

i. A protein has 24 amino-acids in a chain. How many degrees of freedom does this system have?

ii. Name two physics phenomena causing changes in energy as the protein folds.

iii. Is a probabilistic Roadmap algorithm for path planning complete?

iv. Draw examples of an "oscillator" and a "glider" in the game of Life. (need not be exact)

v. Name two features of WordNet that help organize the lexicon for semantic use?

vi. Could we use a small set of mood words to evaluate sentiment of the public via email analysis?

vii. If one could classify situations in chess into winning / not-winning, how would that help us formulate heuristics?

viii. define a transformation grammar that can take an L shape and convert it i) to $<$, and ii) to $>$

ix. Which heuristic for a 24-puzzle will result in a smaller search : out-of-place tiles, or manhattan distance?

x. How many edges need to be matched in a 16x16 edge-matching puzzle?

xi. How does Active Learning reduce the size of a training set?

xii. Where was Sudoku started?

xiii. What is the difference between learning a grammar as in a probabilistic version of E0, vs. learning a grammar as in ADIOS?

xiv. It seems some 4-letter words with P are popular in AI. Give the full forms for [3x2]
a. PCFG b. PHOG c. PLSA