

COMP 5970/6970/6976 Fall 2019 Exam 3

This is a closed-book, closed-notes exam. The only items you are allowed to use are writing implements. Mark each sheet of paper you use with your name and the string “COMP 5970/6970/6976 Exam 3”. If you are caught cheating, you will receive a zero grade for this exam. The max number of points per question is indicated in square brackets after each question. The sum of the max points for all the questions is 29, but note that the max exam score will be capped at 28 (i.e., there is 1 bonus point, but you can’t score more than 100%). This exam consists of 9 multiple-choice questions followed by two open questions. You have exactly 50 minutes to complete this exam. Good luck!

Multiple Choice Questions - write the letter of your choice on your answer paper, NOT on the question sheet which you may keep

1. Hyper-heuristics are particularly well suited for: [2]
 - (a) Sequential EAs
 - (b) Synchronous Parallel EAs
 - (c) Asynchronous Parallel EAs
 - (d) all of the above
 - (e) none of the above
2. Learning Classifier Systems (LCS) improve on classic expert systems by: [2]
 - (a) storing knowledge in the form of rules
 - (b) improving accuracy through reinforcement learning
 - (c) gaining new knowledge by generating rules employing evolutionary computation
 - (d) all of the above
 - (e) a and b
 - (f) a and c
 - (g) b and c
 - (h) none of the above
3. In generative neuroevolution, such as those employing Lindenmayer Systems: [2]
 - (a) Lamarckian evolution is supported because you can employ backpropagation on the neural network weights
 - (b) Lamarckian evolution is supported because Lindenmayer systems accurately reflect how nature evolve complex lifeforms such as trees
 - (c) Lamarckian evolution is not supported because the phenotype-to-fitness mapping is not one-to-one
 - (d) Lamarckian evolution is not supported because the genotype-to-phenotype mapping is highly complex
4. In the context of Assignment 2c, the risk of evaluating a Pac-Man controller against a single Ghost controller is: [2]
 - (a) As this undersamples the opposing population in the most extreme, the Pac-Man controller’s fitness estimation is almost guaranteed not representative of the true value.
 - (b) There may be another Ghost who is better than the Ghost sampled against, yet Pac-Man does better against that Ghost, violating the transitive property of the Ghost fitness function.
 - (c) While it does undersample the opposing population, the greatly reduced evaluation time allows for many more generations to be run, and this is generally a beneficial tradeoff.
 - (d) b and c
 - (e) none of the above

5. A cooperative CoEA is a CoEA where: [2]
- (a) each population tries to solve its own problem without harming the fitness of any of the other populations
 - (b) the populations are symbiotic species
 - (c) each population is a different species representing part of a larger problem
 - (d) all of the above
 - (e) none of the above
6. The Iterated Prisoner's Dilemma is a game where in successive rounds, two individuals without means of communication must decide whether to defect or cooperate; this game: [2]
- (a) is technically not a competitive coevolution problem because it is a single population problem
 - (b) is technically not a competitive coevolution problem because it is a single species problem
 - (c) both of the above
 - (d) none of the above
7. A CIAO plot which is grayish (i.e., predominately either a uniform gray or grainy gray mix of light and dark) is indicative of: [2]
- (a) mediocre stability
 - (b) cycling
 - (c) disengagement
 - (d) answers a and b
 - (e) answers a and c
 - (f) answers b and c
 - (g) answers a, b, and c
 - (h) none of the above
8. To battle the "human bottleneck" in interactive evolution, one can employ: [2]
- (a) surrogate fitness functions
 - (b) small population sizes
 - (c) multi-objective EAs for problems that exhibit a mixture of quantitative and qualitative aspects
 - (d) crowdsourcing
 - (e) all of the above
 - (f) a, b, and d
 - (g) b and d
 - (h) none of the above

9. Increasing primitive granularity in hyper-heuristics is: [2]

- (a) a terrible idea, because we want more sophisticated algorithms, not more primitive
- (b) a sensible idea, because it increases the search space, so the probability increases that a high quality solution is represented
- (c) a problematic idea, because it expands the search space, so the expected run time to find the global optimum increases
- (d) a wonderful idea, because as the granularity approaches Turing completeness, the hyper-heuristic performance will converge on GP optimality
- (e) all of the above
- (f) b and c
- (g) b and d
- (h) b, c, and d
- (i) c and d
- (j) none of the above

Regular Questions

10. How are conflicting rules in the match set of a Learning Classifier System resolved? [3]

11. Given the following bit strings v_1 through v_5 and schema S

$v_1 = (11101110111101) \text{ fitness}(v_1) = 0.3$

$v_2 = (10110010001101) \text{ fitness}(v_2) = 0.1$

$v_3 = (00001010011010) \text{ fitness}(v_3) = 1.0$

$v_4 = (01001110111001) \text{ fitness}(v_4) = 1.9$

$v_5 = (11001011110101) \text{ fitness}(v_5) = 1.7$

$S = (***** ***)$

- (a) Compute the *order* of S . [1]
- (b) Compute the *defining length* of S and show your computation. [2]
- (c) Compute the fitness of S and show your computation. [2]
- (d) Do you expect the number of strings matching S to increase or decrease in subsequent generations? Explain your answer! [3]