COMP 5660/6660 Fall 2026 Exam 2 Key

This is a closed-book, closed-notes exam. The sum of the max points for all the questions is 56, but note that the max exam score will be capped at 52 (i.e., there are 4 bonus points, but you can't score more than 100%). You have exactly 50 minutes to complete this exam. Good luck!

Multiple Choice Questions

1. In Assignment 1b, you were asked to tune your EA to significantly improve its performance. Which of the following turned out to be significant in tuning your EA? [4 pts]

   (a) balancing μ with k in k-tournament parent selection

   (b) choice of parent selection method (e.g., k-tournament, fitness-proportionate selection)

   (c) choice of mutation operator

   (d) mutation rate

   Select one of:

   - a [1]
   - b [1]
   - c [1]
   - d [1]
   - a and b [2]
   - a and c [2]
   - a and d [2]
   - b and c [2]
   - b and d [2]
   - c and d [2]
   - a, b, and c [3]
   - a, b, and d [3]
   - a, c, and d [3]
   - b, c, and d [3]
   - **a, b, c, and d**
   - none of a, b, c, nor d [0]

2. If we employ self-adaptation to control the value of penalty coefficients for an EA with an evaluation function which includes a penalty function, then: [4 pts]

   (a) the penalty coefficients will be self-adapted to cause fitness improvement just like, for instance, mutation step sizes

   (b) cannot be done because it is inherently impossible to self-adapt any part of the evaluation function

   (c) the penalty coefficients will be self-adapted, but the increase in fitness achieved may not be correlated with better performance on the objective function

   Select one of:

   - a [2]
   - b [1]
   - **c**
   - none of a, b, nor c [0]

3. In Assignment 1c, you implemented a constraint satisfaction EA employing a penalty function, which relaxed the strict constraints from Assignment 1b on invalid solutions in order to explore the search space more broadly. Is it reasonable to use smaller values of μ in 1c versus 1b? [4 pts]

   (a) Yes, because the EA can now differentiate between solutions containing invalidities, which should make it easier for evolution to discover better solutions, even in small populations with no valid solutions.

   (b) Yes, because the penalty function implements soft constraints rather than hard constraints, thus increasing the survival rate, which allows the use of smaller populations.

   (c) No, because a lower μ would require a higher penalty coefficient to maintain a constant selective pressure, and the higher penalty coefficient will prematurely lead to almost valid solutions not surviving.

   (d) No, because small populations lack the genetic diversity needed for sufficient penalized individuals to "swim across the sea of invalidity" to find new areas of validity.

   Select one of:

   - **a**
   - b [1]
   - a and b [3]
   - c [0]
   - d [1]
   - c and d [1]
   - none of a, b, c, nor d [1]

4. Which of the following are improvements on the base Evolutionary Programming (EP) algorithm: [4 pts]

   (a) Borrowing the self-adaptation of mutation step sizes from Evolutionary Strategies.

   (b) Self-adapting covariance matrices.

   (c) Combining Gaussian & Cauchy distributions to generate random mutations.

   (d) Tuning the recombination parameters to optimize performance for a specific problem class.

   Select one of:

   - a [1]
   - b [1]
   - c [1]
   - d [0]; one of the defining characteristics of EP is that it considers each trial solution as representing a different species, hence there is no recombination
   - a and b [2]
   - a and c [2]
   - a and d [0]
   - b and c [2]
   - b and d [0]
   - c and d [0]
   - **a, b, and c**
   - a, b, and d [1]
   - a, c, and d [2]
   - b, c, and d [2]
   - a, b, c, and d [3]

- none of a, b, c, nor d [0]

5. In Evolution Strategies with uncorrelated mutation with n step sizes, the conceptual motivation for updating the mutation step sizes with the formula $\sigma_i' = \sigma_i \cdot e^{\tau' \cdot N(0,1) + \tau \cdot N_i(0,1)}$ is: [4 pts]

   (a) the sum of two normally distributed variables is also normally distributed

   (b) the common base mutation $e^{\tau' \cdot N(0,1)}$ allows for an overall change of the mutability, guaranteeing the preservation of all degrees of freedom

   (c) the coordinate-specific $e^{\tau \cdot N_i(0,1)}$ provides the flexibility to use different mutation strategies in different directions

   Select one of:

   - a [1]
   - b [1]
   - c [1]
   - a and b [3]
   - a and c [3]
   - b and c [3]
   - **all of a, b and c**
   - none of a, b, nor c [0]

6. Increasing computational power (compute) affects Koza's A-to-I ratio (or AI ratio, meaning "artificial-to-intelligence" ratio) as follows: [4 pts]

   (a) increases A (numerator)

   (b) decreases A (numerator)

   (c) increases I (denominator)

   (d) decreases I (denominator)

   (e) increases the AI ratio

   (f) decreases the AI ratio

   Select one of:

   - a, c, and e [3]
   - a, c, and f [2]
   - a, d, and e [3]
   - **a and e**
   - b, c, and f [0]
   - b, d, and e [0]
   - b, d, and f [0]
   - b and f [2]
   - c and f [0]
   - d and e [0]
   - none of a, b, c, d, e, nor f [1]

7. In the context of multi-objective problem solving, the term scalarisation refers to combining single objective fitness scores into a weighted cumulative fitness score. This approach suffers from the following drawbacks: [4 pts]

   (a) scalarisation commonly is a computationally expensive process

   (b) the implicit assumption that all the user's preferences can be captured before the range of possible solutions is known

   (c) for repeatedly solving different instances of the same problem, either the user's preferences are assumed to be static, or the user needs to repeatedly provide new weightings

   Select one of:

   - a [0]
   - b [2]
   - c [2]
   - a and b [1]
   - a and c [1]
   - **b and c**
   - all of a, b, and c [3]
   - none of a, b, nor c [0]

8. The current GP practice of strongly limiting the role of mutation in favor of recombination is because: [4 pts]

   (a) Recombination tends to increase genetic diversity in GP, unlike mutation which contrary to standard EAs (that employ a linear representation), has a tendency to destroy critical alleles.

   (b) The generally shared view that in GP, crossover has a large shuffling effect, acting in some sense as a macromutation operator.

   (c) Mutation tends to cause excessive bloat in GP, unlike recombination which has a natural parsimony pressure effect.

   Select one of:

   - a [0]
   - **b**
   - c [0]
   - a and b [2]
   - a and c [0]
   - b and c [2]
   - a, b, and c [1]
   - none of a, b, nor c [0]

9. Rechenberg's 1/5 success rule refers to: [4 pts]

   (a) the minimum successful mutation rate threshold necessary for an Evolution Strategy to reach the global optimum

   (b) the ratio of offspring created by mutations versus recombination in Evolutionary Programming

   (c) a rule of thumb for the optimal ratio of successful versus total mutations in Evolution Strategies where mutation step size is increased if the ratio is greater than 1/5 and decreased if the ratio is smaller than 1/5

   (d) the minimum ratio of successful offspring creation versus total offspring creation in order for a parent to survive to the next generation

   Select one of:

   - a [2]
   - b [0]
   - **c**
   - d [1]
   - none of a, b, c, nor d [0]

10. Which of the following are drawbacks of Evolutionary Programming (EP) not utilizing a recombination operator? [4 pts]

    (a) Without recombination, the algorithm cannot create new alleles and so is stuck with the alleles generated during initialization.

    (b) The algorithm's inability to combine independent local improvements into a single offspring.

    (c) The algorithm is prone to stagnation if mutation step-sizes and strategies aren't carefully calibrated and controlled during evolution.

    Select one of:

    - a [0]
    - b [2]
    - c [2]
    - a and b [1]
    - a and c [1]
    - **b and c**
    - a, b, and c [3]
    - none of a, b, nor c [0]

11. Say you want to purchase a new laptop and care most about maximizing performance and minimizing price. You collect performance data and pricing on ten different laptops and then you normalize both the performance data and the pricing which results in the following table, where higher affordability numbers indicate lower price:

ID	Performance	Affordability
1	1	1
2	10	10
3	2	5
4	5	2
5	3	9
6	4	10
7	7	7
8	5	9
9	10	4
10	9	9

(a) List for each element which elements it dominates; indicate elements with their IDs. [4 pts]

ID	Dominates
1	None
2	1, 3-10
3	1
4	1
5	1,3
6	1,3,5
7	1,3,4
8	1,3,4,5
9	1,4
10	1,3,4,5,7,8

(b) Show the population distributed over non-dominated levels, like some multi-objective EAs employ, after each addition of an element, starting with element 1 and ending with element 10 increasing the element number one at a time; indicate elements with their IDs. So you need to show ten different population distributions, the first one consisting of a single element, and the last one consisting of ten elements. [12 pts]

After adding element 1:

Level 1: 1

After adding element 2:

Level 1: 2

Level 2: 1

After adding element 3:

Level 1: 2

Level 2: 3

Level 3: 1

After adding element 4:

Level 1: 2

Level 2: 3,4

Level 3: 1

After adding element 5:

Level 1: 2

Level 2: 4,5

Level 3: 3

Level 4: 1

After adding element 6:

Level 1: 2
Level 2: 4,6
Level 3: 5
Level 4: 3
Level 5: 1

After adding element 7:

Level 1: 2
Level 2: 6,7
Level 3: 4,5
Level 4: 3
Level 5: 1

After adding element 8:

Level 1: 2
Level 2: 6,7,8
Level 3: 4,5
Level 4: 3
Level 5: 1

After adding element 9:

Level 1: 2
Level 2: 6,7,8,9
Level 3: 4,5
Level 4: 3
Level 5: 1

After adding element 10:

Level 1: 2
Level 2: 6,9,10
Level 3: 7,8
Level 4: 4,5
Level 5: 3
Level 6: 1