NAT-DL: Self-study questions for Revision Set 8 (week 10)

Below you can find example questions from previous exams. It unlikely, but not impossible that one of the questions below re-occurs in this year's exam. You may like to check past NAT papers already before this final tutorial, so that you can also discuss other potential exam questions.

Please consider a few hints for answering exam questions. They may not apply to any exams of any other courses, but they will likely apply to this year's NAT exam.

- Preparation is key. There are two ways to prepare:
  1. Answer a choice of questions from past papers, tutorials, and quizzes
  2. Formulate your own questions based on the course material: A typical question is "How does *X* relate to *Y*?", where *X* and *Y* can be principles, facts, algorithms, applications etc. If your question appears very easy, then you can be pleased to know. If it's off moderate difficulty, you may like to search for an answer in the course materials and beyond. If your question is too hard/to detailled/too open-ended, then it is unlikely to occur in the exam, but you may like to keep it in mind for future research
- Past papers contain exam questions of courses from previous years. The course does not remain the same over time, so sometimes you will see a question related to content that was not even mentioned in this year. If you have made sure that it really wasn't mentioned, then you can rest assured that such a question will not occur in this year's exam.
- During the exam, read each question carefully, and plan your answer. A question can have several parts, and in order to get a maximum number of marks, all parts need to be mention in the answer. It is possible to answer all parts in one paragraph and let the marker sort your text into their marking categories, but it is more safe to answer each part in a separate paragraph, even if there may be a bit redundancy.
- If several marks can be achieved for a single question, then you may need to write down several points/examples/reasons/aspects/properties or whatever the question is asking for. Typically (but not always), you get one mark per item, so you can guess how many you need to list.
- If more than one mark is possible for a more challenging answer, then usually one mark is for the plain answer, one is for a good explanation that this is a good answer, one is for limitations of scope of the answer, one is for generalisation, one is for discussion, … but no worries, if this was expected then carefully read the question which needs to indicate this (see above!).

Questions from past papers:

1. Explain the difference between "roulette wheel" selection and a "tournament" selection in genetic algorithms. For what conditions would you prefer on of these selection mechanisms?
2. The Boolean satisfiability problem is as follows: given a formula in conjunctive normal form, find an assignment of Boolean values to the variables in the formula which make the entire formula true. A formula in conjunctive normal form is a conjunction (and) of a set of clauses, with each clause being a disjunction (or) of atomic variables or negated atomic variables. For example:

(L1 OR L2 OR NOT L3) AND (L3 OR NOT L4 OR NOT L5) AND (NOT L2 OR L4 OR L5)

(i) Describe how you could apply the canonical genetic algorithm to the Boolean satisfiability problem.

(ii) How could you improve upon your canonical GA solution for this problem?

3. Your company has the task of designing controllers for spaceships in the “Killer Asteroid” computer game: asteroids travel across the screen destroying any spaceship they hit, unless the spaceship fires at the asteroid and destroys it first. The spaceships’ aim is to survive for as long as possible. Spaceships may move by firing two thrusters, one on each side of the spaceship; the spaceships are working in a gravity-free environment so firing the thrusters causes a spaceship to translate or rotate or both. A spaceship’s missiles are ejected from its nose, so the spaceship must be pointing in the right direction if a missile is to intersect the path of the asteroid. Assume a spaceship has sensors looking in 8 quadrants around it that can measure quantities like distance to the nearest asteroid in each quadrant, speed of asteroid, bearing of asteroid with respect to the spaceship, etc. Describe how you would use GP to come up with a controller for a spaceship, giving details of and justifying all assumptions you make.
   Discuss the problems your solution might have in evolving good controllers quickly. How would your solution fare if the game is made more complicated, e.g. by requiring the spaceships to avoid each other or allowing them to hide behind planets.
4. Give an example of a deceptive fitness function. How can the deceptiveness of a fitness function be measured if the global optimum is known? What is a fully deceptive fitness function? Explain whether the schema theorem applies to the case of deceptive fitness functions.
5. Formulate the travelling salesperson problem as a combinatorial optimisation problem.
6. The Bigbooks book company want to deliver its books to stores around the country. It hires lorries from Toptrucks and delivers several parcels to each store. Each lorry starts at Bigbooks’s depot, where it is filled with parcels. Bigbooks wishes to hire as few lorries as possible so the space inside them needs to be used efficiently. The company wishes to use as little fuel as possible, so the distance travelled by the lorries should be as short as possible. You may assume that weight is not an issue, i.e. a fully loaded lorry is not too heavy. How would you use ACO to produce routes for the lorries? In your answer be sure to describe the use you make of heuristics. How can local heuristics be employed in ant colony optimisation problems?
7. Describe one crossover and one mutation mechanism that would be suitable for use with real-valued chromosomes. In each case use an example to illustrate your answer.
8. State the schema theorem briefly, using either words or an equation. Why might the schema theorem not give an accurate prediction of a GA’s performance?
9. What characteristics of an optimisation problem do you need to consider when deciding whether to solve it using a GA or GP? What are the advantages and disadvantages of using GP especially when compared to using a GA?
10. How can a genetic algorithm be used to evolve a plan? Give details of all the steps involved. What changes must be made in a genetic algorithm to obtain a steady-state genetic algorithm?
11. How is the diversity of the population maintained in the Differential Evolution (DE) algorithm?

12. Particle swarms optimisation (PSO) has been shown to be applicable to similar problems as GA. How does the representation of TSP by PSO relate to GA? Can PSO be expected to outperform GA for the TSP?
13. Ant colony optimisation (ACO) can be applied to solving time tabling problems. For example, in a railroad company wants to schedule trains such that the maximal number of passengers are transported, delays are minimised, and customer satisfaction is maximised.
    a. Discuss the application of one of the ACO algorithms to this problem. For this purpose describe the problem representation, define a desirability heuristic, name relevant constraints, formulate a rule for pheromone update, and write down the probability rule.
    b. Write down the main steps of ACO in pseudocode.
    c. In Ant Colony Optimisation the pheromone trace is reduced in each step by an amount proportional to the evaporation rate. What is the effect of a high or a low evaporation rate on the population and the quality of solutions? Explain how these effects can be used in order to achieve a good performance of the algorithm.
    d. The algorithm converges to a good solution only after a very long time. What may have led to the slow convergence? How can the convergence speed be improved?
    e. The algorithm does not converge to a good solution even after very long time. How can the performance be improved?
    f. Discuss for the present problem how local search is used to improve solutions.
    g. ACO was inspired by the behaviour of real ants. In what way does the algorithm deviate from the biological original?