TD2: Ant Colony Optimization

Instructions:

- You should do every exercise except the hard exercises marked with (✿).
- Fully read the exercise statement before starting to solve it, and always justify your answers.
- If an exercise asks for an algorithm, describe it in pseudo-code or Python.

Exercise 1: Influence of the number of ants. In this exercise, we will reflect on the influence of the number of ants and iterations on the efficiency of ant colony optimization approaches.

1. What is the difference between running Ant System with m ants for k iterations and running it with $m \cdot k$ ants for 1 iteration? Which option will (probably) perform better?
2. What if we run the second option for 2 iterations instead of one?
3. What about running with 1 ant for $m \cdot k$ iterations?
4. Finally, do the above responses change if we use Max-Min Ant System instead?

Exercise 2: Influence of the probabilistic parameters. Consider now running Ant System using $\alpha = 0$, $\beta = 1$ and a single ant, applied to TSP.

1. Describe the functioning of this algorithm.
2. Compare it now to a simple greedy algorithm, which obtains a tour by choosing the closest unvisited vertex until every vertex is visited.
3. If we now change the value of β in Ant System, how does it change the behavior of the algorithm?

Exercise 3: Exploration or exploitation (Exam 2025). The Max-Min Ant System algorithm has several mechanisms that make it different from basic Ant System:

- Upper and lower limits on the pheromone values;
- Pheromone update only on the best solution found so far (global best);
- Periodic trail reinitialization, which resets the pheromone values.

For each of these, explain if it favors exploration or exploitation, and why.

Exercise 4: Dynamic problems. Ant colony optimization is considered a good approach when dealing with dynamic problems.

1. Can you explain why this is the case, using as example a version of TSP where nodes appear and disappear at random?
2. What considerations would you require in the implementation of an algorithm to the problem above?

Exercise 5: The knapsack problem. In the knapsack problem, we have a bag with limited space, and want to pack items into the bag so as to carry as much value as possible. For each of the n items, we are given its weight w_i and its value v_i, and we are also given the capacity of the bag B (in weight). Our goal is to select a subset of the items such that the total weight of the selected items is at most B, and the total value is maximized.

1. Describe a solution representation and natural interest function for the problem.
2. Propose a method for solution construction that always produces feasible solutions.

Exercise 6: Bi-objective TSP. Bi-objective TSP is a variant of TSP where each edge has both a distance cost and a time cost, and we are interested in obtaining different solutions representing different trade-offs of time to distance.

This exercise concerns the modifications that are necessary to apply ACO approaches to bi-objective TSP (and multi-objective problems in general).

1. While in the past we had a single global best solution, in the multi-objective setting we may have a set of "best" solutions. Can you describe the property of these "best" solutions?
2. Consider a variant of the usual algorithms, where each edge has a single pheromone value (as usual). What changes are needed to adapt these approaches to bi-objective TSP.
3. Next, we consider two colonies of ants with different pheromones, one which considers only distance costs and the other only time costs. How does this change the algorithm and the results?

In multi-objective optimization, the most common version of ant colony optimization uses an improvement of the two-colony variant, where each ant has a parameter λ, and it considers the cost and pheromones as:

$$\eta_{uv} = \lambda \cdot \eta_{uv}^{\text{dist}} + (1 - \lambda) \cdot \eta_{uv}^{\text{time}}$$
$$\tau_{uv} = \lambda \cdot \tau_{uv}^{\text{dist}} + (1 - \lambda) \cdot \tau_{uv}^{\text{time}}$$

4. How can you simulate the two-colony approach by setting the values of λ for each ant?
5. Finally, one possible approach is to take a random value of λ for each ant. What is the advantage of this version relative to the previous ones?