UNIVERSITETET I OSLO

Faculty of Mathematics and Natural Sciences

Exam in:	**INF3490/4490 — Biologically Inspired Computing**
Day of exam:	**December 3rd, 2012**
Exam hours:	**14:30 – 18:30**
This examination paper consists of 3 pages.	
Appendices:	**None**
Permitted materials:	**None**

Make sure that your copy of this examination paper is complete before answering.

Part A

1 (weight 5%)
Compare and contrast constrained optimization and unconstrained optimization.

2 (weight 5%)
What helps simulated annealing get out of local minima/maxima?

3 (weight 5%)
Are long runs beneficial in evolutionary algorithms? Explain your answer.

4 (weight 4%)
Describe the purpose of parent selection and survivor selection in evolutionary algorithms.

What is the general effect of using different implementations of these selection mechanisms?

5 (weight (4% + 5%))
Why is the one-point crossover operator not suitable for use with permutation representation in evolutionary algorithms?

Describe a more suitable crossover operator.

6 (weight 3%)
List three termination conditions used in evolutionary algorithms.

7 (weight 4%)
Describe how *self adaptation* is used in some evolutionary algorithms.

8 (weight 5%)
Describe evaporation and its role in ant colony optimization.

9 (weight 5%)
Is it beneficial to add domain knowledge to evolutionary algorithms? Justify your answer.

Part B

1 (weight 4%)
Can a simple perceptron represent the XOR function? Explain your answer.

2 (weight 6%)
How does a multilayer perceptron need to be adjusted so that it can represent an increasingly complex decision boundary? Briefly describe how the numbers of layers, nodes, and weights need to be changed.

3 (weight 6%)
Describe how overfitting can occur specifically in the cases of multilayer perceptrons and support vector machines.

4 (weight 6%)
Explain the method and purpose of early stopping.

5 (weight 4%)
What is the purpose of feature mapping in support vector machines?

6 (weight 4%)
What assumption are we making about the features in our dataset when we use a naïve Bayes classifier?

7 (weight 2%)
Describe two distance measures used in determining nearest neighbors in the k-NN algorithm.

8 (weight 4%)
What is the effect of using different values of the parameter *k* in the k-NN algorithm?

9 (weight 2%)
Define a suitable state space and an action space for the game of chess in a reinforcement learning context.

10 (weight 4%)
Compare the effect of using each of the following reward functions on learning to play the game of chess in a reinforcement learning context:

i. r = +1 if the agent wins; r = -1 if it loses; and r = 0 otherwise.

ii. Same as i., but additionally r = +1 when the agent captures an opponent's piece.

11 (weight 4%)
Recall the ε-greedy policy whereby an agent chooses a random action with probability ε, and the action with the highest estimated value otherwise. Why do we typically decrease the value of ε over time?

12 (weight 5%)
How do self-organising maps help reduce the dimensionality of data? Why is dimensionality reduction useful?

13 (weight 4%)
What is the difference between supervised and unsupervised learning?

----------End----------