Order of presentations: TBA

Time

Total time 9 minutes: 7 minutes for presentation, 2 minutes for questions.

Please measure your time and make sure you make good use of it. I will have a timer and you will have to stop at the 7 minutes mark. After that I might ask you a question or let somebody from the audience ask a question.

If you have a pdf or PowerPoint presentation, send it to me (even a few minutes before the session starts) so that I can put it on my laptop and we save some time.

The day of the presentation, please have your laptop on and your presentation ready if you want to use your laptop.

Content

Given the short time, I will ask you **NOT** to give an overview of what you will be talking about (list of contents or outline). Besides, all of the presentations should have the same general outline, so there is no need to go over it. With that said, the following is what you should include in your presentation. Be brief but clear. The number of slides is just a suggestion, not a hard requirement:

- Key ideas/contributions (1 slide)
    - What problem are you solving?
    - What approach are you using?
    - E.g. "Applied *Genetic Algorithms* to the Salesman problem", "Used *X* and *Y* learning algorithms on datasets *W* and *Z*", etc.
    - Brief high-level explanation of what the problem (e.g. Salesman problem) or dataset is about. If you worked on the Sudoku, you do NOT need to explain it.
- Approach (2 to 4 slides)
    - Explain the details of how you applied the approach. Do NOT explain in detail the approach (assume the audience knows it). For example, if you worked on Genetic Algorithms (GA), you do not need to explain how a GA works. Or if you used machine learning (ML), you do not need to explain how the algorithms you chose work. If you want to say a few words about it is ok, but concentrate more on explaining how you applied it to your particular problem. So, again with the examples:
        - For GA, explain how you represented your individuals (a string of size *X* where each character represents …), how you generated initial population, how you defined the fitness or evaluation function, how you applied selection, crossover and mutation.
        - For ML, more details of the dataset, what features are there, what size is it (example of input table), and what the prediction/learning task is. Briefly mention what each ML algorithm is supposed to do with this dataset and how.

        - For Sudoku, probably you will need to explain a little bit more about the approaches that you tried and how you applied them.
- Evaluation: experiments and results (2 to 3 slides)
    - Show the experiments you conducted to evaluate your approach. For example:
    - For ML
        - How you measured performance of the algorithms. How you partitioned your dataset into training and testing.
        - Show learning curve graph. Check the Nov. 16 class slides (http://www.csee.umbc.edu/courses/671/fall10b/Class22-MachineLearning.pdf) for those about "Evaluation Methodology".
    - For GA:
        - Show results for different values of the parameters of the GA (number of iterations, population size, prob. of crossover, prob. of mutation). You could also observe execution time with respect to input size.
        - Show graphs with the results. For example, on one graph, you can show the performance of the GA as the number of iterations increases (same for population size).
    - Sudoku
        - Show results for different boards (similar to HW2)