

Algorithms for the web and for social networks

Final evaluation

Paolo Boldi

The final evaluation for the course “Algorithms for the web and for social networks”, part of the Bertinoro International Spring School 2012, will be based on an assignment that the students are supposed to carry out in groups of two or three people (or up to four, for the “Software writing” task). Each group should:

- elect among the possible types of assignments listed below
- decide on which specific topic their essay will focus
- contact the teacher to check the appropriateness of their choice and to establish a submission deadline.

The teacher will evaluate the submitted essay, based on the criteria specified below. The following types of assignments are proposed:

1. Project proposal;
2. Reference review;
3. Software writing.

1 Project proposal

The students are supposed to write a proposal for a two-year post-doc grant, concerning or related to one of the topics covered during the course. The proposal, that cannot exceed *four pages*, must contain (and will be evaluated on the basis of):

- a synthetic, credible description of the type of research that the student will be accomplishing;
- a brief review of the state-of-the-art in the area, with a clear indication of why the proposed research should be considered innovative and promising;
- a series of milestones that the student is supposed to reach during the research period, with expected timings and possibly with an indication of deliverables (reports, software etc.) to be produced for each milestone;

- an indication of the kind of measures that will be considered to testify the quality of the results obtained (with a discussion on the datasets and baselines used).

2 Reference review

The students are involved in the writing of a paper concerning or related to one of the topics covered during the course. (S)he should produce, for the sake of their coauthors, a review of the state-of-the-art in the area. The review can be in the form of a brief essay (not exceeding *three pages*) or in the form of a series of slides for a 45-minute presentation. The students are supposed to attach to his/her essay/presentation a brief piece of text explaining the supposed content of the paper (it is *not* required that the innovative part of the paper be described; it is sufficient to contextualize it within the area). The work will be evaluated on the basis of the following:

- the review must be up-to-date and cover all the main recent work in the area, and give also a historic overview of the most important papers and books on the subject;
- a specific indication of the contribution of every paper should be given, using a common notation and nomenclature that make the papers comparable with one another;
- (if applicable) the review must describe the datasets and baselines typically used to evaluate the research, and an indication of the current state-of-the-art of the performance (accuracy, quality) obtained;
- the reader/listener must have, at the end, a precise idea of the roadmap followed in the research and of the accomplishments obtained so far.

3 Software writing

The students are supposed to add a new functionality or feature to an existing IR- or Web-related open-source software project. The project chosen should be active and well-known, and the feature should be significant and interesting. The students are supposed to provide:

- a brief description (two/three pages) of what the new functionalities are and how they will be integrated in the framework: these specifications must be provided to the teacher *before starting to work*, to allow him to decide about the appropriateness of the task;
- the source code of the new functionalities;
- full easy-to-follow instructions on how the teacher can compile/install the framework and the new functionalities, along with the data and instructions to run a demo.

The evaluation will be based on the following criteria:

- quality of the software produced (comments, documentation, unit tests, coverage tests, etc.);
- how close is the solution implemented to the specifications;
- how likely it is that the new feature will be used by somebody (based on its performance, scalability, usability, usefulness, etc.).