STAT 548 PhD Qualifying Papers Guide and Expectations

Benjamin Bloem-Reddy

September 16, 2020

Choosing a paper

At the end of this document is a list of papers and project ideas that I am interested in supervising as Qualifying Papers (QPs). I may update it throughout the year. I am happy to discuss any other paper that you are interested in and think might be appropriate. I am generally interested in theoretical and methodological aspects of statistics and machine learning, and various problems in applied probability. For example, I have worked on probabilistic modeling and Bayesian inference (including Bayesian nonparametrics) for network data, incorporating probabilistic symmetry into deep learning methods, SMC/MCMC sampling methods, and characterizing limiting behavior (degree properties, distributional limits) of Preferential Attachment models for random graphs.

Expectations

If you are interested in doing a QP with me, the first step is to email me to schedule a one-on-one meeting. Please use the words "Qualifying Paper" in the subject line of your email (I have a filter set up so that I don't miss these). Please come to this meeting prepared to discuss the following:

- Your background.
- Your long-term research interests (it's okay if these are not yet well-defined).
- Why you are interested in the particular paper/project.
- When you will submit your report (typically about four-six weeks after we meet).
- The details of the QP project and report.

Report[1]

The report should have the following structure:

- **Summary (maximum of five pages).** The first section of the report should provide a summary of the paper and the problem(s) it addresses, including previous work, the contributions of the work in question (WIQ) (e.g., novel techniques, problem formulations, or synthesis; important results; subsequent work), *why* the results and contributions are important, and any limitations or shortcomings (e.g., restrictive assumptions or flawed methodology). The aim of this section is for you to synthesize the findings of a body of work and clearly present the important points.
- **Mini-proposals for research projects (maximum of two pages per proposal, at least one proposal—no limit on number of proposals).** Each proposal should describe a research project that applies, extends, generalizes, adapts, or addresses shortcomings of the WIQ.[2] Seemingly unrelated

[1] With thanks to Trevor Campbell, from whom I am borrowing liberally.

[2] It should go without saying that you can't propose the project you work on as part of the QP report.

ideas inspired by the WIQ are also fine. A proposal should concisely describe: the primary problem to be addressed; an approach (or multiple approaches) for addressing the problem; any technical or conceptual sub-problems; the potential impact of the project. You are not expected to pursue any of these projects (though we can talk more if you would like to). The aim of this section is to get you thinking creatively about research, and to begin developing the skills necessary for writing research proposals.

- **Paper-specific project results (no page limit).** Each potential QP listed below has a brief description of a related project. We will discuss the project in detail in our initial meeting, and we can meet again (as many times as necessary) before the report due date. Your grade *will not* be affected by how good the results look, whether your approach improves on past work, or whether you achieve the initial goal of the project. I will use this project to evaluate your research potential, which includes:
  - clearly formulating a research question;
  - setting up a useful mathematical framework for the problem;
  - thinking creatively and independently to develop a solution;
  - relating the problem to existing work, in other fields if necessary;
  - being resourceful and asking questions when necessary;
  - learning from and moving past the inevitable setbacks;
  - reformulating the research problem when necessary;
  - implementing new methods in code (when applicable);
  - choosing appropriate experiments and metrics;
  - communicating and reflecting on progress, setbacks, and results;
  - thinking of future research directions.

  The report should be submitted as a GitHub repository downloaded from the template here: https://github.com/ben-br/qp-template/. The template includes a LaTeX style file that should be used for the report. (Detailed instructions for usage can be found in the repository's README file.)

  Any experimental/numerical results should be reproducible. All code should be reusable, clearly commented/documented, and exist in a GitHub repository to which I have access as a collaborator. I prefer Julia or Python (in that order), but we can talk about using other languages if necessary.

Resources

- Some resources on technical/mathematical writing:
  - Nancy Heckman's page: http://ugrad.stat.ubc.ca/~nancy/writing/
  - Harry Joe on mathematical writing and typesetting in LaTeX: https://www.stat.ubc.ca/~harry/papers/
  - Trevor Campbell's "How to Explain Things" talk
  - Knuth, Larrabee, and Roberts on mathematical writing: http://www.jmlr.org/reviewing-papers/knuth_mathematical_writing.pdf
  - Halmos on writing mathematics: https://www.math.uh.edu/~tomforde/Books/Halmos-How-To-Write.pdf (a transcribed, searchable PDF with some typos: https://entropiesschool.sciencesconf.org/data/How_to_Write_Mathematics.pdf)
- Getting started with Git: chapters 1 and 2 of https://git-scm.com/book/en/v2 should be all you need for this report.

Benjamin Bloem-Reddy
benbr@stat.ubc.ca

Papers and projects (last updated: September 16, 2020)

1. MCMC for Hawkes and related point process models

Paper: J. G. Rasmussen. "Bayesian Inference for Hawkes Processes". *Methodology and Computing in Applied Probability* 15.3 (2013), pp. 623–642

Categories: MCMC, point processes

Implement the MCMC algorithms of this paper, and one of the following:

Option 1. extend to include edge effects ([MR05] may be useful); or

Option 2. quantify the error in the posterior resulting from neglecting edge effects ([MR06] may be useful); or

Option 3. modify the algorithms for models that include inhibitory terms.

2. CNNs without a group

Paper: R. Kondor and S. Trivedi. "On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups". *Proceedings of the 35th International Conference on Machine Learning*. Ed. by J. Dy and A. Krause. Vol. 80. 2018, pp. 2747–2755. URL: http://proceedings.mlr.press/v80/kondor18a.html

Categories: theory, methodology, machine learning

Explore generalizations of this paper to sets of transformations that do not form a group, e.g., monoid, semigroup, other sets of transformations with algebraic structure useful to the problem. (You may find a probabilistic approach [BT20] useful.)

References

[BT20] B. Bloem-Reddy and. Teh. "Probabilistic Symmetries and Invariant Neural Networks". *Journal of Machine Learning Research* 21.90 (2020), pp. 1–61. URL: http://jmlr.org/papers/v21/19-322.html.

[KT18] R. Kondor and S. Trivedi. "On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups". *Proceedings of the 35th International Conference on Machine Learning*. Ed. by J. Dy and A. Krause. Vol. 80. 2018, pp. 2747–2755. URL: http://proceedings.mlr.press/v80/kondor18a.html.

[MR05] J. Møller and J. G. Rasmussen. "Perfect simulation of Hawkes processes". *Advances in Applied Probability* 37.3 (2005), pp. 629–646.

[MR06] J. Møller and J. G. Rasmussen. "Approximate Simulation of Hawkes Processes". *Methodology and Computing in Applied Probability* 8.1 (2006), pp. 53–64.

[Ras13] J. G. Rasmussen. "Bayesian Inference for Hawkes Processes". *Methodology and Computing in Applied Probability* 15.3 (2013), pp. 623–642.