Aditya Gangrade

Email: adityagangrade92@gmail.com, agangra2@andrew.cmu.edu; Website: adityagangrade.wordpress.com

Research Interests

Broadly: Statistical Learning Theory; Online Learning; Information Theory; Statistics;
Concretely: Reliable and resource-efficient learning; High-dimensional and nonparametric testing

Education

Ph.D. and *M.S.* in Systems Engineering, Boston University (2022)
Advisors: Bobak Nazer, Venkatesh Saligrama.
B.Tech. (with honours) in Electrical Engineering, IIT Bombay (2014)

Academic Positions

Postdoctoral Researcher, Department of Statistics and Data Science, Carnegie Mellon University (2/22-present)
Graduate Student Researcher, Division of Systems Engineering, Boston University (9/15-1/22)
Research Intern, School of Technology and Computer Science, Tata Institute of Fundamental Research (2/15-8/15)
Junior Research Fellow, Dept. of Electrical Communication Engineering, Indian Institute of Science (6/14-1/15)

Research Projects

Safe Bandits

- Studied enforcement of round-wise stochastic safety constraints in bandit problems.
- Demonstrated ineffectiveness of prior aggregate-constraint strategies for round-wise settings.
- Developed 'doubly-optimistic' schemes for Multi-Armed and Linear bandits as an alternative to prior pessimistic schemes. Showed that doubly-optimistic schemes have favourable regret and safety properties without assumptions.

Testing Log-Concavity

- Studied non-parametric testing for log-concavity of a distribution in batch and sequential setting.
- The batch setting yielded the first provably valid and consistent test for log-concavity.
- Showed that test martingales are powerless in the sequential setting, and developed powerful E-process based tests.

Selective Classification

- Investigated learning paradigms where the learner can abstain from making predictions on queries.
- Characterised selective classification as a primitive for inference time resource-accuracy tradeoffs.
- Constructed a novel formulation, that yields new insights on the complexity of the problem.
- Developed efficient and statistically effective methods for the problem.
- Developed foundations for SC with limited feedback in the online learning framework.
- Studied applications of these ideas to hybrid cloud-edge computation.

Knowledge Distillation

- Designed a novel and principled method for knowledge distillation that is cognisant of the mismatch of model capacities between students and teachers.
- Structurally, the method deemphasises dataset examples that are 'hard' for the teacher in an adaptive manner, thus spending student capacity in an efficient way.
- The method is flexible and can augment any KD loss, and leads to marked improvement in student accuracy, particularly in the low student capacity regime.

Structural Testing of Networks and Graphical Models

- Investigated the hypothesis testing of the underlying structure of network-based models.
- Yielded new mathematical insights into the question of when structural testing is easier than the corresponding recovery problem in Ising and Stochastic Block Models.
- Established, a phase transition in the complexity of structural testing driven by the size of the changes; Constructed statistically and computationally efficient schemes to detect changes.

Nonparametric Regression Methods

- Constructed a nonparametric regression scheme over the class of piecewise-linear delta-convex functions.
- The resulting method is highly expressive, has strong statistical rates, and is efficiently implementable.

Teaching Experience

As a graduate student

- TA'd two upper level mechanical engineering courses, both focusing on control theory. These are
  - Vision, Robotics, and Planning (*Spring '20; BU; 700 level*)
  - Atmospheric Flight Mechanics and Control (*Spring '17; BU; 400 level*)
- Ran labs, occasional recitations for the latter. Mainly grading work for the former.

As an undergraduate student

- Extensive experience as a TA for introductory classes in Physics and Mathematics.
  - Electricity & Magnetism (*Autumn '11, Summer '12, Spring '13,'14; IITB; 100 level*)
  - Quantum Mechanics (*Spring '13; IITB; 300 level*)
  - Calculus (single and multiple variable) (*Autumn '13; IITB; 100 level*)
  - Calculus (intensive program for entrants) (*Autumn '12; IITB; 100 level*)
- Typical responsibilities included running weekly tutorials for ~ 40 students, that entailed reviewing theory covered in lectures, and solving example problems in detail.
- I was coordinating TA in the Spring '13 and '14 runs of E. &M., and also involved with designing evaluation.
- The I.P.E. was an experimental program as part of the affirmative action strategy of IITB, which entailed giving extensive instruction in English & Mathematics to admitted students with underpriviliged backgrounds that lacked sufficient grounding in the same. My work consisted of about 3 hours of in class lecturing and problem solving per week, with a focus on development of concepts, as well as of the students' confidence with academic and non-academic aspects of the institute.

Publications & Preprints

Chen, T., A. Gangrade, and V. Saligrama (2022a). "A Doubly Optimistic Strategy for Safe Linear Bandits". In: *arXiv preprint arXiv:2209.13694*.

Chen, T., A. Gangrade, and V. Saligrama (2022b). "Strategies for Safe Multi-Armed Bandits with Logarithmic Regret and Risk". In: *International Conference on Machine Learning* (ICML).

Dunn, R., A. Gangrade, L. Wasserman, and A. Ramdas (2021). "Universal Inference Meets Random Projections: a Scalable Test for Log-Concavity". In: *arXiv preprint arXiv:2111.09254*.

Gangrade, A., A. Kag, A. Cutkosky, and V. Saligrama (2021). "Online Selective Classification with Limited Feedback". In: *Advances in Neural Information Processing Systems* (NeurIPS). Spotlight presentation.

Gangrade, A., A. Kag, and V. Saligrama (2021). "Selective Classification via One-Sided Prediction". In: *International Conference on Artificial Intelligence and Statistics* (AISTATS).

Acar, D. A. E., A. Gangrade, and V. Saligrama (2020). "Budget Learning via Bracketing". In: *International Conference on Artificial Intelligence and Statistics* (AISTATS).

Gangrade, A., B. Nazer, and V. Saligrama (2020). "Limits on Testing Structural Changes in Ising Models". In: *Advances in Neural Information Processing Systems* (NeurIPS).

Siahkamari, A., A. Gangrade, B. Kulis, and V. Saligrama (2020). "Piecewise Linear Regression via a Difference of Convex Functions". In: *International Conference on Machine Learning* (ICML).

Gangrade, A., P. Venkatesh, B. Nazer, and V. Saligrama (2019). "Efficient Near-Optimal Testing of Community Changes in Balanced Stochastic Block Models". In: *Advances in Neural Information Processing Systems* (NeurIPS).

Gangrade, A., B. Nazer, and V. Saligrama (2018). "Two-Sample Testing can be as Hard as Structure Learning in Ising Models: Minimax Lower Bounds". In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 6931–6935.

Gangrade, A., B. Nazer, and V. Saligrama (2017). "Lower bounds for two-sample structural change detection in Ising and Gaussian models". In: *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, pp. 1016–1025.

Miscellaneous

BU Division of Systems Engineering Dissertation Award (2022)

Best Paper Award at BU Centre for Information and Systems Engineering's Graduate Student Workshop for work on piecewise linear regression (2021)

Awarded the *Dean's Fellowship* by the College of Engineering, BU (2015-16)

Reviewer for IEEE Trans. Inf. Th., NeurIPS, COLT, AISTATS, ICLR. (2017-)

Outstanding Reviewer at ICLR '21 and NeurIPS '21