Research Statement of Yisong Yue

Objective

Making sense of digital information is a growing problem in almost every domain, ranging from scientists needing to stay current with new research, to companies aiming to provide the best service for their customers. This issue represents a central aspect of the so-called "Big Data" phenomenon.

What is common to many Big Data problems is not only the scale of the data, but also the complexity of the processes that generated the data. Most importantly, much of the data is the result of human interaction, and reflects the full complexity involved in human decision making. Examples include purchase data, search logs, communication transcripts, and GPS traces. To effectively use such data, it is often not sufficient to merely scale up existing machine learning methods. One must also understand and model human decision processes in order to accurately learn from observable human behavior (cf. [2]). Furthermore, although datasets have grown dramatically in recent years, human attention has not. As a result, meaningful consumption of large datasets often requires tailoring to the requirements of specific applications and individual people.

Building upon this insight that interesting Big Data problems often pertain to people and how they interact with a digital environment, my research agenda is centered around developing ***principled machine learning approaches*** that not only efficiently utilize the abundant amount of available raw data, but also ***accurately interpret*** observable human behavior, ***intelligently interact*** with humans, and more generally learn with ***humans in the loop***.[1] Developing effective approaches requires a confluence of disciplines including machine learning, decision theory, natural language processing, computational psychology, and the social sciences.

In my research methodology, I aim to ***identify motivating applications*** that can be enabled or improved by new machine learning approaches, ***formulate new learning problems*** that characterize the key technical challenges, ***develop efficient algorithms*** with learning-theoretic guarantees, and ***validate the effectiveness*** of new methods in practical settings. As a consequence, my research profile spans a broad spectrum from theoretical advances all the way to live usability studies (cf. [10, 11]), and demonstrates the end-to-end value of my research contributions.

Prior Work

My prior work has already demonstrated both the richness of research questions arising from this area, as well as the end-to-end research contributions for a number of applications.

Interactive Structured Prediction for Recommender Systems. Virtually all information needs are ambiguous or multi-faceted, and virtually all datasets contain abundant amounts of redundant information. Information systems that do not account for these factors will quickly drown users with an overwhelming amount of redundant, or worse, irrelevant information. My dissertation

[1] See also `http://tinyurl.com/c2lkw8e`

work proposed the ***first principled learning approach*** for training structured coverage models to recommend diverse content that "covers" all relevant facets or topics [12, 4]. My postdoctoral work on Linear Submodular Bandits [10] shows how to interactively personalize such models to the specific tastes of individual users. I proposed an interactive learning framework that builds upon and unifies bandit theory[2] and structured coverage models, and developed a ***provably near-optimal algorithm*** with ***efficient regret convergence guarantees***. In a live user study, we found that ***90% of users preferred our approach*** over competing recommender algorithms.

Online Learning from Preference Feedback. One major issue in learning from user interactions is that most interactions reveal only implicit feedback. Even something seemingly explicit such as a star rating requires calibration across users and items. One reliable approach is to infer preferences (is A better than B?) rather than absolute feedback (how good is A?) [5, 2]. My dissertation work proposed a novel framework, called the Dueling Bandits Problem, for analyzing interactive algorithms that learn from preference feedback [13, 7, 8], and developed the ***first provably efficient no-regret algorithms*** for this setting. I further developed extensions to accommodate violations in preference transitivity (which are often observed in practice), and can result in ***orders-of-magnitude improvement*** in performance [14]. In collaboration with researchers at Yahoo!, I also developed contextual approaches to improve the statistical power of preference elicitations in search applications, resulting in a ***40% reduction*** in the amount of feedback required [9, 2].

Other Work. The technical principles underlying my research have applications beyond user-centric information systems. As part of my postdoctoral work at the iLab at Carnegie Mellon University, I applied these principles to tackle the problem of automated decision support for emergency medical services [15].[3] Despite the differences in application, there is nonetheless a similar notion of multi-facted coverage to optimize: adaptively positioning an entire fleet of ambulances to best "cover" the expected distribution of future emergency requests. Building upon this intuition, I developed a ***provably near-optimal approach*** for ambulance fleet allocation and dynamic redeployment that can provide a ***70% reduction*** in the number of unserviced emergencies.

Future Directions

As shown in my prior work, I am interested in developing ***end-to-end machine learning approaches*** that are motivated by ***real applications***, yield ***new theoretical advances***, and demonstrate ***tangible practical impact***. A central theme of my research focuses on developing interactive approaches that conduct goal-oriented experimentation for eliciting targeted human feedback while causing minimal usability interference. Such approaches can better interpret observable behavior and improve system utility for a variety of applications, ranging from helping domain experts model complex phenomena to personalizing recommender systems for large populations of users.

Structured Active Learning. When human experts employ machine learning to help them analyze their domains of study, they often go through multiple rounds of refining before arriving at a suitable model. Such settings typically benefit from active learning approaches that aim to elicit the most informative feedback for refinement. Furthermore, the complexity of modern datasets are

[2]The multi-armed bandit problem is a framework used for analyzing tradeoffs between exploration and exploitation.
[3]See also `http://www.cmu.edu/news/stories/archives/2012/july/july24_improvingemergencyservice.html`

often best understood using rich structured models that both capture complex statistical regularities and facilitate ease of inspection. ***One example is my work*** on structured sentiment classification [6], where the model justifies its predictions by extracting the best supporting sentences or phrases. Moving forward, I am interested in developing structured active learning approaches that can simultaneously leverage the representational power of complex structured models while provably retaining the efficiency of conventional active learning methods for unstructured models.

Furthermore, to facilitate the modeling and understanding of increasingly complex datasets, I am interested in developing structured interactive approaches that both model the data and elicit refinement at multiple levels of granularity or abstraction. Effective approaches should be able to characterize and optimally balance the relative tradeoffs of refinement elicitation not only between different levels of abstraction but also for different predictions within the same level.

Learning from Rich Interactions. Many applications involve inherently complex or difficult sensemaking tasks that cannot be adequately solved by simply predicting a static set of results. This problem is further exacerbated in settings where observable human behavior contains insufficient information for the model to make accurate predictions. To address these issues, I am interested in developing machine learning approaches that holistically integrate with richer interfaces that provide enriched communication channels between humans and their digital collections. ***One preliminary example is my work*** on Dynamic Ranked Retrieval [1], which uses a decision-theoretic model to dynamically generate search results given real-time feedback. Compared to static rankings, our approach yields ***provably superior utility***, demonstrates significantly improved performance for queries with ambiguous intent, and exhibits a ***20% improvement*** in overall performance.

One example of a rich interface is the Apolo system [3], which allows users to seamlessly annotate and cluster items to facilitate goal-oriented knowledge discovery, and was used to effectively solve a literature review task that was difficult to solve using conventional systems. To effectively model such interfaces, we must identify a new collection of interaction "primitives" that can characterize these new enriched communication channels between the human and the system, and formulate new learning problems to reason about how to learn using these primitives. We must also develop new algorithms to address the computational challenges that will inevitably arise from modeling, learning, and predicting using rich interactions.

Collaborative Learning Systems. Modern information systems do not service individual users in a vacuum, but rather must provide service simultaneously for large populations of users. Effective and broadly applicable learning approaches should have both the flexibility to model and personalize to individual users, as well as the ability to intelligently balance the exploration/exploitation tradeoff for entire populations of users. ***One preliminary example is my work*** on bandit transfer learning [11], which simultaneously addresses both issues to a limited degree with ***provably efficient convergence guarantees***. In a live user study, we found that our approach could ***more efficiently personalize to 80% of users*** without sacrificing model flexibility.

More generally, I am interested in extending collaborative learning approaches to accommodate the aforementioned models of complex sensemaking tasks and rich interfaces. Effective approaches would require methods that can jointly learn both the transient sensemaking goals within a single rich interactive session as well as the long-term interests that persist within populations of users.

References

[1] Christina Brandt, Thorsten Joachims, Yisong Yue, and Jacob Bank. Dynamic ranked retrieval. In *ACM Conference on Web Search and Data Mining (WSDM)*, 2011.

[2] Olivier Chapelle, Thorsten Joachims, Filip Radlinski, and Yisong Yue. Large scale validation and analysis of interleaved search evaluation. *ACM Transactions on Information Science (TOIS)*, 30(1):6:1–6:41, 2012.

[3] Duen Horng Chau, Aniket Kittur, Jason I. Hong, and Christos Faloutsos. Apolo: Making sense of large network data by combining rich user interaction and machine learning. In *ACM Conference on Human Factors in Computing Systems (CHI)*, 2011.

[4] Thorsten Joachims, Thomas Hofmann, Yisong Yue, and Chun-Nam Yu. Predicting structured objects with support vector machines. *Communications of the ACM (CACM), Research Highlight*, 11(52):97–104, November 2009.

[5] Filip Radlinski, Madhu Kurup, and Thorsten Joachims. How does clickthrough data reflect retrieval quality? In *ACM Conference on Information and Knowledge Management (CIKM)*, 2008.

[6] Ainur Yessenalina, Yisong Yue, and Claire Cardie. Multi-level structured models for document-level sentiment classification. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2010.

[7] Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. In *Conference on Learning Theory (COLT)*, 2009.

[8] Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. *Journal of Computer and System Sciences (JCSS), Special Issue on Learning Theory*, 2012. doi:10.1016/j.jcss.2011.12.028.

[9] Yisong Yue, Yue Gao, Olivier Chapelle, Ya Zhang, and Thorsten Joachims. Learning more powerful test statistics for click-based retrieval evaluation. In *ACM Conference on Information Retrieval (SIGIR)*, 2010.

[10] Yisong Yue and Carlos Guestrin. Linear submodular bandits and their application to diversified retrieval. In *Neural Information Processing Systems (NIPS)*, 2011.

[11] Yisong Yue, Sue Ann Hong, and Carlos Guestrin. Hierarchical exploration for accelerating contextual bandits. In *International Conference on Machine Learning (ICML)*, 2012.

[12] Yisong Yue and Thorsten Joachims. Predicting diverse subsets using structural svms. In *International Conference on Machine Learning (ICML)*, 2008.

[13] Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *International Conference on Machine Learning (ICML)*, 2009.

[14] Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *International Conference on Machine Learning (ICML)*, 2011.

[15] Yisong Yue, Lavanya Marla, and Ramayya Krishnan. An efficient simulation-based approach to ambulance fleet allocation and dynamic redeployment. In *AAAI National Conference on Artificial Intelligence (AAAI)*, 2012.