Research Statement

Haoqi Zhang

January 19, 2012

My research interests focus on advancing our ability to design and study social and economic systems on the Web to promote desired participant behaviors and outcomes. I call this *computational environment design*. My research broadly spans and draws from the fields of artificial intelligence (AI), human-computer interaction (HCI), decision science, and the emerging discipline of economics and computer science (EconCS). I am particularly interested in problems motivated by the transformative conditions on the Internet today and promised by the continuing digital revolution, notably the ability to (a) recruit individuals with diverse expertise to join problem solving efforts and discussions via crowdsourcing and social media, (b) track complex individual or group behaviors over time across multiple interactions, and (c) experiment with alternative designs and get rapid feedback by observing user behavior. These conditions introduce new opportunities to intelligently harness the availability of the crowd, and to leverage the ability to track behavior to gain new insights into participants' motivations and decision-making processes, based on which to improve designs. Some particular questions of interest include:

- How can we effectively recruit a *crowd* and coordinate its actions to tackle complex tasks?
- How can we *automatically design environments* by systematically discovering interventions tailored to the preferences and capabilities of participants despite imperfect information?
- How do system-wide policies and rules influence user behavior in online communities?

These questions present new challenges in multi-agent interaction for both HCI and EconCS, e.g., in providing effective interfaces or workflows for collaboration among individuals, and in properly handling incentives and designing rules of interaction. AI's role is multi-faceted: such questions suggest new opportunities in how computers can take part during problem solving, and orthogonally, in how computers can help us understand the connection between design and behavior. My work in this space has been to create applications that illustrate novel paradigms and design patterns, develop general methods and theories that identify new principles and techniques, and conduct empirical studies that through measurement offer insights into how the design of actual systems influence user behaviors and outcomes. These considerations surface throughout my research, as highlighted below.

Human Computation and Crowdsourcing

Crowdsourcing Complex Tasks

Over the last decade, *human computation* has established itself as a powerful paradigm for incorporating human intelligence in problem solving efforts in which machines cannot yet tackle the problem alone. While most existing human computation systems involve simple tasks that are easy to parallelize across individuals, there has been considerable recent interest in *crowdsourcing complex tasks* that require more sophisticated coordination and optimization. While efforts have focused primary on designing *human computation algorithms or workflows* that decompose a task into more manageable subtasks that the crowd tackles independently, such approaches are less applicable to complex tasks that involve a set of global constraints. In examples such as writing and itinerary planning, good, feasible solutions rely on the composition as a whole, and are marked by interdependence among solution components.

Using itinerary planning as a case study [10], I designed and built a system called Mobi, which takes as input a natural language mission describing one's wishes and requirements for a day trip, and generates as output a full itinerary satisfying that mission. Mobi takes inspiration from *groupware*, and illustrates a novel *crowdware* paradigm that provides a *workspace* in which a crowd of individuals contribute opportunistically based on their knowledge and expertise and the current solution context, and in which the system (indirectly) coordinates the crowd problem solving effort. To promote effective contributions, Mobi displays automatically generated reminders that alert participants of violated constraints and suggests actions for resolving such constraints. Experiments show that all else equal, displaying such reminders significantly increases the rate at which the crowd resolves violated constraints. Furthermore, in an end-to-end user study in which subjects entered personal missions based on their intended travels, all subjects found that the crowd-generated itineraries satisfied their stated requirements, and would be useful in real life.

For both crowdware and workflow paradigms, I see opportunities to leverage the crowd not only at the endpoints of computation, but also in guiding the control flow of an algorithm, and even in generating plans that define the problem-solving process [9]. This has led to novel applications such as CrowdPlan [4], a system that allows a user to enter a high-level mission as a search query (e.g., 'I need to manage my inbox better,' or 'I want to live a healthier life'), and uses a crowd to generate a plan that consists of concrete steps one can take to accomplish the mission and associated web resources for tackling each step. Results from user studies show that CrowdPlan is able to provide users with diverse, actionable steps and useful roadmaps for accomplishing their missions.

In other efforts, I demonstrate how existing human computation design patterns can be effectively combined to create powerful new design patterns and applications in which the crowd's performance matches that of experts. One example is a novel dual-pathway structure for iterative tasks [5], that induces participants to improve upon previous solutions without requiring any external form of quality control such as voting. The structure is particularly apt for problems in which checking the quality of solutions is costly, as is the case for audio transcription. In an experiment with 147 Harvard undergraduates, we found that transcripts generated through the dual-pathway structure approached the accuracy of professional transcription, and were significantly more accurate than the best transcripts in a parallel condition. Another example is PlateMate [6], which takes as input a photograph of a plate of food and organizes a crowd to produce estimates of nutritional information (e.g., calories, fats, proteins, carbohydrates) as output. PlateMate incorporates both our observations of a trained nutritionist's work process and appropriate control of the crowd, and in user studies provided estimates that matched the accuracy of trained nutritionists.

Task Routing

Engaging a crowd to tackle complex tasks relies not only on effective coordination, but on recruiting individuals with appropriate expertise to join the problem solving effort and providing proper incentives for inducing contributions. In social networks, an individual's knowledge often extends beyond their own expertise on tasks and topics to knowledge about the expertise of others. I am interested in principles and methods for *task routing*, that aim to harness the ability of people or automated agents to both contribute to a solution, and to route tasks or subtasks to others who they believe can also effectively solve and route [9].

In work on task routing for prediction tasks [8], I introduce *routing scoring rules* that via appropriate rewards bring truthful contribution of information and optimal routing of the task into a Perfect Bayesian Equilibrium under common knowledge about the amount of information available to each person on the network. But while the common knowledge assumption may be reasonable for small organizations, it is unlikely to hold for large social networks where each person's information about the competencies of others is limited to a local neighborhood. To address this challenge, I introduce a family of *local routing rules* that isolate simple routing decisions in equilibrium that can be easily computed using shared local knowledge, while still promoting effective routing and information aggregation. In doing so, *I design incentive schemes that explicitly enable equilibrium behavior for which the inference required of participants is tractable.*

Automatic Environment Design

When designing crowdware and incentive schemes for task routing, I reason explicitly about the participants' capabilities and decision-making processes, based on which I construct effective designs. In the meanwhile, I have also been working on methods for enabling *automatic environment design*, that take as input a set of available interventions, the objective of the designer, and a model of the interaction between environment, participants, and behaviors, and generate as output an intervention that promotes actions and outcomes meeting the objective whenever such interventions exist.

A key challenge in automatic design is that models are imperfect and incomplete; we have limited knowledge of people's preferences and decision-making processes, and this private information is difficult to elicit directly. But in online settings where the designer can track individual and group behaviors, such information can be indirectly inferred from observing people's decisions over repeated interactions. Coupled with the ability to experiment with alternative designs and receive rapid feedback, this suggests opportunities to construct procedures that automatically drive an objective-oriented, iterative design process by interweaving indirect learning of model parameters with optimization for effective interventions based on the current model.

In work on *policy teaching* [12], I consider a specific environment design problem in a Markov Decision Process setting, in which an interested party associates limited rewards with world states to induce a planning agent to follow a desired policy. The interested party does not know the agent's reward function, but can interact with the agent over time and observe its decisions. I introduce an *active, indirect elicitation* method, that in each round applies an incentive intervention based on a hypothesis on the agent's reward function that is consistent with observations. Based on this method, I devise an algorithm that elicits the desired policy in number of rounds logarithmic in the size of the unknown reward space. For *value-based policy teaching* [11] problems in which the goal is to have the agent follow the best inducible policy for the interested party (which depends on the unknown rewards of the agent), I provide an extension of the elicitation method and associated elicitation heuristics. While the logarithmic convergence bound does not apply to this more complex setting, in simulations I find that we can often elicit the best policy after few rounds.

In follow-up work [7], I present a general domain-independent model of the environment design problem and an associated framework for applying active, indirect elicitation approaches. By framing the problem of finding effective interventions as a learning problem in unknown parameter space, we are able to apply learning methods that are independent of the actual type of intervention considered, and leverage heuristics and techniques from direct elicitation approaches to such indirect problems.

In another direction [1], I relax the planning assumption and consider a setting in which a designer aims to affect the decisions of an adaptive agent whose values for different actions may update with experience. I consider a multi-armed bandit model, and establish strong positive and negative results on the online performance of incentive schemes under different budget settings. An interesting finding is that even with limited knowledge of the agent's adaptive process, an interested party can use observations of the agent's decisions to learn and track his state online, based on which to provide incentives. This result suggests that indirect learning and optimization may be effective even in settings where interventions can influence an agent's future decisions and learning trajectory.

While most of my work on automatic environment design has been theoretical, I am interested in applying automatic design techniques to actual problems on the Web. More recently, I have been pursuing applications to automatically design crowdsourcing tasks and workflows. In work on automatic task design [2], I show that by experimenting with different designs of an image tagging task, we can learn accurate models that predict how design parameters affect the quality and rate of work. Results from experiments show that at the same pay rate, optimized designs provide significantly more high quality tags than baseline designs.

In current work, I am applying automatic design approaches to constructing crowdsourcing workflows. Take sorting as an example: by interweaving learning about how well workers can perform various subtasks (e.g., choosing the pivot, deciding whether something is almost sorted) with optimization over feasible sorting workflows built on such modules, a sorting algorithm can iteratively evolve to tailor itself to the crowd's abilities and the particular sorting task.

Economics of Private Communities

For both automatic and manual approaches to designing social and economic systems on the Web, understanding how participants may respond to system conditions and rules of interaction is crucial. Such considerations are particularly salient in online communities, where policies and social norms shape the patterns of user interaction and the underlying population of users. In this space, I have been studying *private BitTorrent communities*. Private communities built on the BitTorrent protocol develop their own policies and mechanisms for motivating members to share content, in particular by requiring members to maintain a minimum ratio between uploads and downloads. In doing so, private communities effectively establish credit systems, and with them economies.

In a study on DIME [3], a community for sharing live concert recordings, I have through measurements of user and torrent traces over a half year period and interviews with site moderators been able to gain a deeper understanding of the economic forces affecting users in such communities. A key observation is that while the download of files is priced only according to the size of the file, the rate of return for seeding new files is significantly greater than for seeding old files. This effect skews the everyday consumption of files toward newer files, and makes it difficult for infrequent and low bandwidth users to earn credit. Interestingly, the DIME community's decision to require a lower than usual share ratio still allows these 'less fortunate' users to participate, which illustrates how the design of policies reflect social norms, and affect the composition of users and behaviors accordingly.

Research Directions

I am excited to pursue a broad agenda on computational environment design that spans theoretical, empirical, and experimental styles. I highlight a few directions of interest below.

In crowdsourcing complex tasks, I am interested in developing other crowdware systems along with theoretical models, in order to fundamentally understand the spectrum between crowdware and workflow paradigms. I am particularly interested in challenges of scale, and in developing approaches that would allow a large number of participants to effectively contribute to complex problems that are hard to decompose.

A key challenge is that as the number of ideas and the size of the solution grows, it becomes difficult if not impossible for any given individual to keep track of the entire solution context. One idea for overcoming this challenge is to apply the concept of a 'semantic zoom' to create *task platforms* that generalize both crowdware and workflow paradigms, where by presenting context at the right level of detail, individuals can make effective local contributions while being aware of the effect of their actions on the global solution. From the workflow perspective, task platforms are algorithms that determine and display at any given time a set of available subtasks, and for any selection construct an interface that provides the necessary context and functionality for that subtask. From the crowdware perspective, task platforms consist of multiple workspaces that cover different but interdependent aspects of the problem. I am interested in exploring both perspectives, and in developing frameworks, methods, and applications that leverage this concept.

I am also interested in continuing to develop general principles and methods for task routing, that aim to effectively and efficiently harness the diverse expertise of participants in a system. In addition to exploring numerous task-level issues, I am interested in organizational-level issues related to distributing streams of tasks in a manner that takes into account people's solving and routing abilities over a spectrum of tasks, as well as people's changing levels of attention, motivation, and availability, and the corresponding need for balancing the load across participants.

Finally, I am interested in applying and developing automatic environment design in a wide range of real world scenarios. In the near term, I am particularly interested in studying how one can automatically design webpages to promote desired usage patterns, e.g., that lead to increases in levels of contribution, comprehension, and awareness. In the longer term, I am interested in pursuing opportunities for automatic design in the physical world, where advances in ubiquitous sensing and the increasing digitization of real world spaces have the potential to enable interactions in which spaces can automatically configure themselves to promote desired behaviors and outcomes.

In pursuing my research, I have actively sought out collaborations with researchers with backgrounds in AI, HCI, EconCS, and Theory, whose expertise and diverse perspectives have not only helped to advance particular projects but also my personal development as a researcher and thinker. In this year alone, I have published papers with fourteen different collaborators on seven separate projects. I look forward to the prospect of forming my own research group, and mentoring both graduate and undergraduate students to pursue interesting, creative work.

References

[1] Y. Chen, J. Kung, D. C. Parkes, A. D. Procaccia, and H. Zhang. Incentive design for adaptive agents. In *AAMAS '11*, 2011.

[2] E. Huang, H. Zhang, D. C. Parkes, K. Z. Gajos, and Y. Chen. Toward automatic task design: a progress report. In *HCOMP '10*, 2010.

[3] I. A. Kash, J. K. Lai, H. Zhang, and A. Zohar. Economics of BitTorrent communities. *In submission to WWW '12*, 2012.

[4] E. Law and H. Zhang. Towards large-scale collaborative planning: Answering high-level search queries using human computation. In *AAAI '11*, 2011.

[5] B. Liem, H. Zhang, and Y. Chen. An iterative dual pathway structure for speech-to-text transcription. In *HCOMP '11*, 2011.

[6] J. Noronha, E. Hysen, H. Zhang, and K. Z. Gajos. Platemate: Crowdsourcing nutrition analysis from food photographs. In *UIST '11*, 2011.

[7] H. Zhang, Y. Chen, and D. C. Parkes. A general approach to environment design with one agent. In *IJCAI '09*, 2009.

[8] H. Zhang, E. Horvitz, Y. Chen, and D. C. Parkes. Task routing for prediction tasks. To appear in *AAMAS '12*, 2012.

[9] H. Zhang, E. Horvitz, R. C. Miller, and D. C. Parkes. Crowdsourcing general computation. In *the CHI workshop on crowdsourcing and human computation*, 2011.

[10] H. Zhang, E. Law, R. C. Miller, K. Z. Gajos, D. C. Parkes, and E. Horvitz. Human computation tasks with global constraints. To appear in *CHI '12*, 2012.

[11] H. Zhang and D. C. Parkes. Value-based policy teaching with active indirect elicitation. In *AAAI '08*, 2008.

[12] H. Zhang, D. C. Parkes, and Y. Chen. Policy teaching through reward function learning. In *EC '09*, 2009.