

Research Statement

Shih-Fen CHENG

School of Computing and Information Systems, Singapore Management University

Tel: (65) 6828-0526; Email: sfcheng@smu.edu.sg

19 December 2025

Background

Most of my research efforts are in agent-based models, decentralized optimization, and urban management. In these areas, I enjoy applied research, which includes building systems and working with real-world datasets. My theoretical work is usually done later to explain principles involved in system building and to provide meta-level insights on empirical datasets. Speaking very generally, my primary objectives across all areas are to understand the mechanics of the studied systems and to design and implement novel mechanisms that will improve the performance of these systems. My approach to research is thus naturally interdisciplinary, and I enjoy working with professionals from different application domains.

I use two dimensions in defining my research interests: 1) the methodology dimension, where I am primarily interested in “agent-based modeling and simulation” and “large-scale optimization”; and 2) the application dimension, which includes “commodity market” and various aspects of “urban management”.

Research Areas

All topics I worked on are related to the understanding, modeling, and optimization of complex engineering systems. In the following section, I describe in detail my research contributions in methodological and application fronts. In particular, I will focus on my work on agent-based modeling and simulation, large-scale optimization, commodity market, and urban management.

1. Agent-Base Modeling and Simulation

Agent-based Modeling and Simulation (ABMS) refers to a large body of work that focuses on studying complex system phenomena from the bottom-up. Instead of trying to explain the phenomenon in interest analytically, agent-based approaches focus on finding appropriate individual models and let them interact following real-world-like rules. The purpose is not just to replicate the phenomenon, but also to generate dynamics that can be observed and studied. The ABMS approach is particularly powerful when we have a complex system that cannot be described analytically and when individual's behaviors are quantifiable.

My first exposure to the agent-based methodology started from my graduate school days, during which we have organized and participated in the Trading Agent Competition (TAC). TAC refers a collection of agent-based trading platforms to encourage the study of autonomous strategic reasoning in specific settings (e.g., travel planning, supply chain management, and electricity markets). Since my arrival at SMU, my interests in this area have evolved from focusing only on autonomous “agents” to systems where both software and human agents play important roles. Over my years at SMU, I have made significant

efforts in building ABMS systems in three such areas: 1) commodity markets, 2) taxi fleet operations, and 3) visitor behavior simulations. In each of these three areas, a new idea was explored, and together they form the first important pillar of my research identity.

My first contribution was SimCommodity, which was used in the area of commodity trading. Working with professional traders at the International Trading Institute (ITI) at SMU, and funded by the International Enterprise (part of trade organization within Singapore government), I have developed a unique commodity trading simulation for training students interested in becoming traders; this platform was also later used in ITI's professional programs.

There are two major innovations in this trading simulation platform. The first innovation is its event-driven architecture, which allows an instructor to flexibly define the kind of market scenarios he would want students to experience. These events serve two purposes: 1) it contains human readable contents, which are to be revealed to human traders when the right triggering conditions are satisfied (e.g., time, price), 2) it is used as a trigger to engage autonomous market agents, who then generate appropriate trades according to their roles (e.g., hedgers or speculators), and collectively they produce the resulting market dynamics.

Another innovation is the use of realistic market mechanism in the simulation. The implication of this design is that the scenario designer cannot dictate what the market price must be, and trades made by same-minded human participants can affect market prices in an unexpected manner just as in the real markets. The design principle of SimCommodity has been published both as a research paper [1] and as a software demo [2] at reputable AI conferences. Besides its extensive usage in education, this platform later inspires our research into agent-based approach for regulatory policy analysis.

My second contribution is TaxiSim, which is used for simulating taxi fleet operations. This work is funded by the Future Urban Mobility (FUM) Center under the Singapore-MIT Alliance for Research and Technology (SMART) and the most significant contribution made in this work is the use of big data in quantifying agent's (taxi driver) behaviors. Our data partner in this research is a major taxi fleet operator in Singapore (with more than 50% of the market share in fleet size), who continuously supplied three years (2009–2012) of detail operational data for our research efforts. This dataset contains minute-by-minute GPS coordinates and status of 15,000 taxis and all made trips (origins, destinations, and fares), and the total size of the dataset is well over 1TB (around 35GB per month). By analyzing this dataset, we are able to quantify how real taxi drivers make driving decisions using the cognitive hierarchy model [3] [4]. Our understanding of driver behaviors is then incorporated into a large-scale ABMS capable of simulating more than 10,000 heterogeneous agents in a realistic traffic network (extracted from the OpenStreetMap) [5].

My third contribution is SimLeisure, which is used for supporting our research efforts in dynamic user experience management in the leisure industry. Our primary concern in this area is to best utilize the limited amount of available time of a visitor to maximize his received values from visiting a selected set of "attractions". Unlike the previous two domains, we don't have direct inputs from human participants and we also don't have any reliable historical human traces we can rely on.

To address the issue of lack of data, we introduce the use of "mobile ground surveys" for collecting ground truth. In the mobile ground surveys we conducted, we distribute mobile

App we designed to automatically collect mobility traces from participants and also important decisions made by them (e.g., the decision to queue up for an attraction). Although the percentage of visitors being tracked is extremely low (around 1% in most cases), we manage to demonstrate by simulation that such mobility ground surveys are far superior than other alternatives [6].

The ABMS domain has been, and will remain a backbone of my research identity. I have experimented with different ways of creating and using various ABMS systems, through direct user engagement, big data analytics, and controlled ground surveys. Many of these developed methodologies and systems still play central roles in our ongoing research projects, and these systems are critical in our research efforts in these areas. I am in the process of formalizing the methodologies and principles in designing ABMS platforms based on the above three projects.

2. Large-Scale Optimization

Handling large-scale computation and decentralization has been a major concern in all my research areas. To support the investigation in these areas, I also maintain a separate stream of research focusing on developing suitable methodologies and algorithms for these tasks. There are two major types of methodologies that I am working on:

- The *market-based approaches* that can deal with problems that are inherently distributed. These problems are usually composed of self-interested individuals (i.e., agents) who possess local information and cannot be controlled centrally. Proper incentives and mechanisms need to be introduced in these cases in order to implement any control policy. My methodological contributions in this area are mainly in agent strategy design, game-theoretic analytics, performance analysis, and mechanism design.
- The *decentralized algorithms* that can approximately solve high-dimensional hard optimization problems via parallel computation. Roughly speaking, such techniques will identify manageable sub-problems within a grand problem and regulate how these sub-problems communicate or coordinate with each other (e.g., via game-theoretic framework or Lagrangian framework). The parallel computing is then applied to solve large number of individual sub-problems. Examples of such approach include my published work in coordinated traffic signal control [7] and also multi-tier production system planning [8]. We also experimented Lagrangian relaxation in coordinating agents in a capacitated resource network [9].

The ability to efficiently solve large-scale optimization problems under imperfect conditions (uncertain or even partially observable data, decentralized information structure and control) will continue to be an important concern in all areas of my research. My objective in this area will be to support my other domain-specific interests.

3. Agent-based Commodity Market Analysis

The SimCommodity ABMS platform I created for educational purposes turns out to have been more impactful. Together with an industry collaborator, we recognized the similarity of income distribution and price dynamics between our simulation and the real-world market. This motivates me to further enhance the SimCommodity platform for policy analysis, particularly for investigating the effectiveness of regulatory policies under extreme market conditions. Since most theoretical models break down under extreme conditions, an agent-based model is one of the few frameworks that might work.

One of my recently completed research articles [10] investigates the potential causes for market instability within commodity markets, and we subsequently test whether the widely debated regulatory measure, limiting speculative position limits, can be an effective policy. The key research findings show that the homogeneity among market participants is what causes extreme price movements, and positions limits cannot contain the induced volatility. The only feasible measure seems to be liquidity injection.

This work was well received among practitioners, and shortly after its completion, the Institute for Financial Markets (IFM) awarded us a research grant to analyze the Flash Crash of May 6, 2010. In this project, we adopted an agent-based approach in replicating the underlying market structure that was believed to cause the 2010 Flash Crash. In particular, we replicated the inter-asset connection and fragmented exchanges, and demonstrated that markets can be forced into unbalances and flash crashes when selling pressure mounts and exchanges are implementing circuit-breakers without coordination [11]. This work has important implications to the regulators everywhere, since we demonstrated that the fragmentation of market exchanges (a much less studied topic compared to the fragmentation of markets) is enough for generating flash crashes, and the inter-connectivity of different asset classes will propagate and amplify the extreme price movements from one asset to other related assets. Besides academic publishing and presentations, these two pieces of work have been presented in a number of high-level regulatory meetings, and stimulated many interests and inquiries.

I believe that the use of agent-based technologies would complement classical research methods in finance very well, particularly in extreme conditions where most classical assumptions fail. My current ongoing interest in the area is on the quantification of trader's behaviors and expert's knowledge in commodity trading.

4. Urban Transportation

My interest in urban transportation started in traffic control, where I looked at the coordinated control of traffic signals in a densely populated urban area. The problem is viewed as a combinatorial optimization problem, and I proposed an efficient distributed algorithm for solving it [7]. More recently, I also have begun investigating how the quality of service (QoS) of a taxi fleet can be improved.

In many cities, taking taxis is considered to be the public transportation mode that is closest to owning a car since it allows point-to-point movement. Many factors affect the efficiency of taxi service, for example, government policy (quota management in particular), fare structure, and revenue model. In my research, I focus on investigating taxi drivers' reactions to incentives when the fare is distance-based and drivers pay a fixed rent plus fuel costs and get to keep all earned revenues.

By utilizing the same large-scale dataset used in building TaxiSim, I have analyzed taxi drivers' queueing behaviors at hotspots (e.g., airport, popular tourist attractions, or major hotels) [12] and also their roaming strategies. Despite the difference in circumstances, the conclusion is universal: human decisions are reasonable, but limited when compared to computer generated strategies. The most interesting fact we discover is the strong positive correlation between a driver's "reasoning depth" (derived using cognitive hierarchy model) and his earned revenue. This strongly justifies the need for a decision support system that is adaptive to both real-time demand and supply situations [3], [4].

My most important ongoing research effort in this area is the design and development of mobility-on-demand (MOD) system. In particular, I have focused on approaches that would improve existing taxi fleet's performance and the use of MOD services in last-mile (LM) transportation needs. For the former focus, I have led a team at the Fujitsu-SMU Urban Computing and Engineering Corp Lab to develop a taxi Driver Guidance System (DGS) [18], which incorporates a specially catered real-time data feed from the Land Transport Authority of Singapore. The data feed allows real-time taxi status and location to be sensed, from which we then derived predicted supply and demand in the different part of Singapore and generate optimized recommendations individually for all guided drivers. We have developed and deployed the DGS to test with close to 500 drivers in 2017. With more than one year of testing in 2018, we demonstrate that by following DGS recommendations, drivers can reduce their empty cruising time by almost 30% [19]. We also recently found that providing drivers with real-time information at congested locations such as airport helps reduce their waiting time by 10% and raises daily earnings by \$3.7 [20].

On the stream related to the LM transport, we have proposed to utilize spare taxis (or any available commercial or private vehicles) and coordinate them to offer ride-sharing services. Many important research questions need to be addressed to make this feasible, such as: 1) How should demand be clustered and assigned to taxis? 2) How should travel cost be split? 3) When should we strategically reserve available taxis for future demands? With the funding support from SMART-FUM, we have addressed and published some of these issues [13]. Some major working papers are in the pipeline and will be finished soon to answer these questions.

I am also working with my SMU colleagues on a new established initiatives related to sustainable urban logistics. This is a brand new effort and has yet to produce significant research output, but I do expect this topic to form another important pipeline in my research portfolio.

5. Urban Crowd Coordination

In this area, I look at problems where individuals are moving in a confined space with loose movement/time constraints. Specifically, I am currently working on the following two instances: 1) visitor experience management at leisure destinations (such as theme parks or museums), where the length of visit and the area they can visit are both constrained, yet everything else stays open, and 2) mobile crowdsourcing, where an individual is interested in moving between his preferred origin and destination, following certain route, yet is flexible in making limited detour during the trip to accomplish certain paid tasks.

My research on the problem of visitor experience management at leisure destinations is supported by the Living Analytics Research Center at SMU, and my focus is on the effective use of user guidance so that extreme congestions can be avoided. In particular, my interests in this domain are on crowd management mechanisms, models of visitors, and agent-based simulation (SimLeisure). The most distinct features of my research efforts in this area are the incorporation of "mobile ground surveys" [6] and the consideration of multiple agents [14]. A number of issues related to congestion control have been studied, and are in preparation for publication.

Mobile crowdsourcing is an emerging mode of crowdsourcing scheme, which requires "mobile workers" to physically visit a place and perform some location-specific tasks (photo taking in most cases). There are some successful platforms in existence, yet their

assignment models are myopic and far from optimal. With research support from the Xerox Research India, we are currently working on designing the next-generation mobile crowdsourcing platform. Our ultimate goal in this area is to propose a system that is capable of sensing individual's mobility patterns, and by aggregating all workers' probably patterns and status, generate robust and satisfactory task allocation plans for all participants. To achieve this, our short-term research objectives include: 1) sense individual mobility patterns and contexts from collected traces, and 2) generate robust and close-to-optimal task assignment efficiently. Our initial result on this topic has been published in HCOMP-14, a premier AI conference focusing on human computation and crowdsourcing [15], and a US Patent Application based on our initial research also was filed in conjunction with Xerox Research. Based on this foundation, we further developed the model to handle worker's stochastic daily patterns [16], and have implemented and tested it in a large campus-scale field experiment [17]. Recently, we have also explored how platform operator can utilize task display policy to further improve cost effectiveness [21].

Looking forward, a promising research direction in this space is to explore other creative way of utilizing gig worker population in various other urban services.

Selected Publications and Outputs

- [1] S.-F. Cheng and Y. P. Lim, "An agent-based commodity trading simulation," Twenty-First Annual Conference on Innovative Applications of Artificial Intelligence, 2009, pp. 72–78.
- [2] S.-F. Cheng, Y. P. Lim, and C.-C. Liu, "An agent-based commodity trading simulation," Eighth International Conference on Autonomous Agents and Multiagent Systems, 2009.
- [3] P. Varakantham, S.-F. Cheng, G. J. Gordon, and A. Ahmed, "Decision support for agent populations in uncertain and congested environments," Twenty-Sixth AAAI Conference on Artificial Intelligence, 2012, pp. 1471–1477.
- [4] A. Ahmed, P. Varakantham, and S.-F. Cheng, "Uncertain congestion games with assorted human agent populations," Twenty-Eighth Conference on Uncertainty in Artificial Intelligence, 2012, pp. 44–53.
- [5] S.-F. Cheng and T. D. Nguyen, "TaxiSim: A multiagent simulation platform for evaluating taxi fleet operations," 2011 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2011, pp. 14–21.
- [6] S.-F. Cheng, L. Lin, J. Du, H. C. Lau, and P. Varakantham, "An agent-based simulation approach to experience management in theme parks," Winter Simulation Conference, 2013, pp. 1527–1538.
- [7] S.-F. Cheng, M. A. Epelman, and R. L. Smith, "CoSIGN: A parallel algorithm for coordinated traffic signal control," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 4, pp. 551–564, 2006.
- [8] A. Ghate, S.-F. Cheng, S. Baumert, D. Reaume, D. Sharma, and R. L. Smith, "Sampled fictitious play for multi-action stochastic dynamic programs," *IIE Transactions*, vol. 46, no. 7, pp. 742–756, 2014.
- [9] G. J. Gordon, P. Varakantham, W. Yeoh, H. C. Lau, A. S. Aravamudhan, and S.-F. Cheng, "Lagrangian relaxation for large-scale multi-agent planning," 2012

IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2012, vol. 2, pp. 494–501.

[10] B. Lee, S.-F. Cheng, and A. Koh, “An analysis of extreme price shocks and illiquidity among systematic trend followers,” *Review of Futures Markets*, vol. 18, no. 4, pp. 385–419, 2010.

[11] B. Lee, S.-F. Cheng, and A. Koh, “Would price limits have made any difference to the ‘Flash Crash’ on May 6, 2010,” *Review of Futures Markets*, vol. 19, Special IFM Issue, pp. 55–93, 2011.

[12] S.-F. Cheng and X. Qu, “A service choice model for optimizing taxi service delivery,” Twelfth IEEE Conference on Intelligent Transportation Systems, 2009, pp. 1–6.

[13] S.-F. Cheng, D. T. Nguyen, and H. C. Lau, “Mechanisms for arranging ride sharing and fare splitting for last-mile travel demands,” Thirteenth International Conference on Autonomous Agents and Multi-Agent Systems, 2014, pp. 1505–1506.

[14] C. Chen, S.-F. Cheng, and H. C. Lau, “Multi-agent orienteering problem with time-dependent capacity constraints,” *Web Intelligence and Agent Systems*, vol. 12, pp. 347–358, 2014.

[15] C. Chen, S.-F. Cheng, A. Gunawan, A. Misra, K. Dasgupta, and D. Chander. “TRACCS: A framework for trajectory-aware coordinated urban crowd-sourcing,” Second AAAI Conference on Human Computation and Crowdsourcing, pages 30-40, Pittsburgh, USA, November 2014.

[16] C. Chen, S.-F. Cheng, H. C. Lau, and A. Misra. “Towards city-scale mobile crowdsourcing: Task recommendations under trajectory uncertainties,” Twenty-Fourth International Joint Conference on Artificial Intelligence, pages 1113–1119, Beunos Aires, Argentina, July 2015.

[17] T. Kandappu, A. Misra, S.-F. Cheng, N. Jaiman, R. Tandriansyah, C. Chen, H. C. Lau, D. Chander, and K. Dasgupta. “Campus-scale mobile crowd-tasking: Deployment and behavioral insights,” Nineteenth ACM Conference on Computer-Supported Cooperative Work and Social Computing, to appear, San Francisco, CA, USA, February 2016.

[18] S. S. Jha, S.-F. Cheng, M. Lowalekar, N. Wong, R. Rajendram, T. T. Khiem, P. Varakantham, T. T. Nghia and F. Rahman, “Upping the game of taxi driving in the age of Uber,” in *Thirtieth Annual Conference on Innovative Applications of Artificial Intelligence*, New Orleans, LA, USA, February 2018.

[19] S.-F. Cheng, S. S. Jha, and R. Rajendram. “Taxis strike back: A field trial of the driver guidance system,” *Seventeenth International Conference on Autonomous Agents and Multiagent Systems*, pages 577-584, Stockholm, Sweden, July 2018.

[20] S. Agarwal, S.-F. Cheng, J. Keppo, L. Wang, and Y. Yang, “Information Provision and Search Frictions: Evidence from the Taxi Industry in Singapore,” *The Review of Economics and Statistics*, to appear, 2025.

[21] A. Arslan, F. Kilci, S.-F. Cheng, and A. Misra, “Choice-based Crowdshipping for Next-day Delivery Services: A Dynamic Task Display Problem,” *European Journal of Operational Research*, to appear, 2025.