Yang Liu

Curriculum Vitae

(734) 546-7375
yangl@seas.harvard.edu
www.yliuu.com

ACADEMIC POSITION

Harvard University, Cambridge, U.S.A.
Postdoctoral fellow in SEAS & CRCS (advisor: Yiling Chen) *01/2016 - present*
Visiting scholar (host: Yiling Chen) *11/2015 - 12/2015*

EDUCATION

University of Michigan, Ann Arbor, U.S.A.
Ph.D. in EECS (advisor: Mingyan Liu) *08/2010 - 11/2015*
M.Sc. in Applied Mathematics *08/2012 - 04/2014*
M.Sc. in EECS *08/2010 - 04/2012*

Shanghai Jiao Tong University, Shanghai, China
B.Sc. in Information Security *09/2006 - 06/2010*
B.Sc. (dual degree) in Finance *09/2007 - 06/2010*

University of Hong Kong, Hong Kong, China
Exchange student in Computer Science *08/2008 - 12/2008*

RESEARCH INTEREST

Artificial Intelligence (AI), Machine Learning (ML), algorithmic economics, crowdsourcing, fairness in AI/ML, security/privacy for human data.

PUBLICATION

Yang Liu and Chien-Ju Ho. Incentivizing High Quality User Contributions: New Arm Generation in Bandit Learning. *In the 32nd AAAI Conference on Artificial Intelligence (AAAI), 2018.*

Zehong Hu, **Yang Liu**. Yitao Liang and Jie Zhang. A Reinforcement Learning Framework for Eliciting High Quality Information. *In Machine Learning in the Presence of Strategic Behavior at NIPS, 2017.*

Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalya Mandal and David Parkes. Fair Experimentation: A Calibrated Bandit Framework. *In the Fourth Fairness, Accountability, and Transparency in Machine Learning (FATML)*; *also in the 2017 Conference on Digital Experimentation (CODE@MIT), 2017.*

Yang Liu and Yiling Chen. Machine Learning aided Peer Prediction. *In the 17th ACM Conference on Economics and Computation (EC), 2017.*

Yang Liu. Fair Optimal Stopping Policy for Matching with Mediator. *In the 33rd Conference on Uncertainty in Artificial Intelligence (UAI), 2017.*

Yang Liu and Mingyan Liu. Crowd Learning: Improving Online Decision Making Using Crowdsourced Data. *In the 26th International Joint Conference on Artificial Intelligence (IJCAI), 2017.*

Yang Liu and Yiling Chen. Sequential Peer Prediction: Learning to Elicit Effort using Posted Prices. *In the 31st AAAI Conference on Artificial Intelligence (AAAI), 2017.*

Ji Liu, **Yang Liu**, Angelia Nedich and Tamer Başar. An Approach to Distributed Paramet-

ric Learning with Streaming Data. *In the 56th IEEE Conference on Decision and Control (CDC), invited paper, 2017.*

Yang Liu and Yiling Chen. A Bandit Framework for Strategic Regression. *In the 30th Annual Conference on Neural Information Processing Systems (NIPS), 2016.*

Yang Liu and Yining Wang. Doubly Active Learning: when Active Learning meets Active Crowdsourcing. *In Crowdsourcing and Machine Learning at NIPS, 2016.*

Yang Liu and Yiling Chen. Learning to Incentivize: Eliciting Effort via Output Agreement. *In the 25th International Joint Conference on Artificial Intelligence (IJCAI), 2016.*

Yang Liu and Mingyan Liu. Finding One's Best Crowd: Online Prediction By Exploiting Source Similarity. *In the 30th AAAI Conference on Artificial Intelligence (AAAI), 2016.*

Wenwu He, James T. Kwok, Ji Zhu and **Yang Liu**. A Note on the Unification of Adaptive Online Learning. *In Trans. on Neural Networks and Learning Systems (TNNLS), 2016.*

Yang Liu and Mingyan Liu. An Online Learning Approach to Improving the Quality of Crowd-Sourcing. *In ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), 2015.*
Journal version: ACM/IEEE Transaction on Networking (ToN), 2017.

Yang Liu, Armin Sarabi, Jing Zhang, Parinaz Ardabili, Manish Karir, Michael Bailey and Mingyan Liu. Cloudy with a Chance of Breach: Forecasting Cyber Security Incidents. *In the 24th USENIX Security Symposium (USENIX SEC), 2015.*
Media coverage: Wall Street Journal interview. FICO acquisition.
Patent: No. 62/026,349, 2016

Armin Sarabi, Parinaz Naghizadeh, **Yang Liu** and Mingyan Liu. Prioritizing Security Spending: A Quantitative Analysis of Risk Distributions for Different Business Profiles. *In the 2015 Workshop on the Economics of Information Security (WEIS), 2015.*
Journal version: Journal of Cybersecurity, 2016.

Shang-Pin Sheng, **Yang Liu** and Mingyan Liu. A Regulated Oligopoly Multi-Market Model for Trading Smart Data. *In Smart Data Pricing at INFOCOM, 2015.*

Yang Liu and Mingyan Liu. An Online Approach to Dynamic Channel Access and Transmission Scheduling. *In the 16th International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), 2015.*
Journal version: Springer Handbook of Cognitive Radio, 2017.

Yang Liu and Mingyan Liu. Detecting Hidden Cliques From Noisy Observations. *In the 40th IEEE International Conf. on Acoustics, Speech and Signal Processing (ICASSP), 2015.*

Yang Liu and Mingyan Liu. Detecting Hidden Propagation Structure and Its Application to Analyzing Phishing. *In the 2014 International Conference on Data Science and Advanced Analytics (DSAA), Best Application Paper Award, 2014.*

Yang Liu, Mingyan Liu and Sahand Haji Ali Ahmad. Sufficient Conditions on the Optimality of Myopic Sensing in Opportunistic Channel Access: A Unifying Framework. *In IEEE Transaction on Information Theory (TIT), 2014.*

Yang Liu and Mingyan Liu. To Stay Or To Switch: Multiuser Dynamic Channel Access.

In the 32nd IEEE International Conf. on Computer Communications (INFOCOM), 2013. Journal version: IEEE Transaction on Mobile Computing (TMC), 2014.

Yang Liu, Mingyan Liu and Jing Deng. Evaluating Opportunistic Multi-channel MAC: Is Diversity Gain Worth the Pain? *In the 31st IEEE International Conference on Computer Communications (INFOCOM), 2012.*
Journal version: IEEE Journal on Selected Areas in Communications (JSAC), 2013.
Best Poster Award of 2011 EGS, University of Michigan

WORKING PAPERS

Yang Liu and Yiling Chen. Surrogate Scoring Rules and a Dominant Truth Serum for Information Elicitation. *In submission, 2017.*

Yang Liu, Juntao Wang and Yiling Chen. Wagering Mechanism for Risk-Taking Agents. *In preparation.*

Christos Dimitrakakis, **Yang Liu**, David Parkes and Goran Radanovic. Bayesian Fairness. *In submission, arXiv:1706.00119, 2017.*

Yang Liu, Ji Liu and Tamer Başar. Gossip Gradient Descent. *In submission, 2017.*

Shuran Zheng, Bo Waggoner, **Yang Liu** and Yiling Chen. Active Information Acquisition for Linear Optimization. *arXiv:1709.10061, 2017.*

INVITED TALKS

Surrogate Scoring Rules and Dominant Truth Serum, UPenn Theory Seminar	*10/2017*
Machine Learning aided Peer Prediction, BayesianCrowd	*07/2017*
Machine Learning aided Incentive Design, Dartmouth College CS Colloquium	*05/2017*
Learning and Harnessing the Power of Human Computation, USC ISI	*02/2017*
A Bandit Framework for Strategic Regression, UIUC CSL seminar	*02/2017*
Bandit in Crowdsourcing, Models and Algorithms for Crowds and Networks	*09/2016*
Improving Prediction By Exploiting Source Similarity, ITA	*02/2015*
Anatomy of Network-Level Malicious Activities, statistics seminar, UMich	*05/2014*
Diversity Gains in Multiuser Wireless Network, Qualcomm research	*08/2013*

TEACHING

Graduate Student Instructor of EECS 489 (Computer Networks), UMich *Winter 2012*

AWARDS

Department nomination for Rackham Predoctoral Fellowship	*2015*
Selected for ITA Graduation Day Talk	*2015*
Conference Student Travel Grants • NIPS, USENIX SEC, MobiHoc, SIGMETRICS, INFOCOM	*2013 - 2016*
Best Application Paper Award, DSAA	*2014*
Towner Price (outstanding Ph.D research) Finalist, University of Michigan	*2014*
Rackham Student Travel Grant, University of Michigan	*2012 - 2015*
Best Poster Award (First place) of EGS, University of Michigan	*2011*
Dean's fellowship, EECS:Systems, University of Michigan	*2010 - 2012*
China Undergraduate Mathematical Contest in Modeling, Second Prize	*2009*
Li & Fung Fellowship for exchange student, University of Hong Kong	*2008*
Scholarships acknowledging excellent academic record, SJTU	*2007 - 2009*
National First Prize in Mathematics Olympic Games	*2006*

ACADEMIC SERVICE

Program Committee
Adversarial Reasoning in Multi-Agent Systems at AAMAS 2017.

Conference reviewer
STOC, MLStrat@NIPS, IJCAI, NIPS, ACM EC, ACM SIGMETRICS, IEEE ICC.

Journal reviewer
ACM Trans. on Economics and Computation, Journal of Artificial Intelligence Research, IEEE Trans. on Information Forensics & Security, Trans. on Networking, IEEE TPDS, IEEE Trans. on Information Theory, IEEE TCom, IEEE TMC, JSAC, Neurocomputing,.

INTERNSHIP

Research intern, Qualcomm Research, San Diego, California *Summer 2013*

REFERENCES

Dr. Yiling Chen
Gordon McKay Professor of Computer Science, Harvard University
Maxwell Dworkin 233, 33 Oxford Street, Cambridge, MA 02138
Email: yiling@seas.harvard.edu
Phone: (+1 617) 495-3298

Dr. Mingyan Liu
Professor of EECS, University of Michigan, Ann Arbor
4427 EECS, 1301 Beal Ave, Ann Arbor, MI 48109-2122
Email: mingyan@umich.edu
Phone: (+1 734) 764-9546

Dr. David C. Parkes
George F. Colony Professor of Computer Science, Harvard University
Co-Director, Harvard Data Science Initiative
Maxwell Dworkin 229, 33 Oxford Street, Cambridge, MA 02138
Email: parkes@eecs.harvard.edu
Phone: (+1 617) 384-8130

Dr. Tamer Başar
Swanlund Endowed Chair & CAS Professor of ECE, University of Illinois Urbana-Champaign
Director, Center for Advanced Study
Professor, Coordinated Science Laboratory & Information Trust Institute
1308 W Main Street, Urbana, IL 61801-2307
Email: basar1@illinois.edu
Phone: (+1 217) 979-1283

Dr. Boi Faltings
Professor in the Faculty of Information and Communication Sciences, EPFL
Director, Artificial Intelligence Laboratory
I&C, Station 14, Swiss Federal Institute of Technology, IN, CH-1015 Lausanne (Switzerland)
Email: Boi.Faltings@epfl.ch
Phone: (+41 21) 693-2738/5

Dr. Vijay G Subramanian
Associate Professor of EECS, University of Michigan, Ann Arbor
4112 EECS, 1301 Beal Ave, Ann Arbor, MI 48109-2122
Email: vgsubram@umich.edu
Phone: (+1 734) 615-1915