Jeremy Zhen Yang

Massachusetts Institute of Technology
Sloan School of Management
100 Main Street, E62-586
Cambridge, MA 02142

Phone: (617) 386-6412
Email: yangzhen@mit.edu
Tweet: @jeremyzyang

Education

Massachusetts Institute of Technology (Cambridge, MA)
Ph.D. Management Science (Quantitative Marketing), 2021 (expected)
Advisor: Sinan Aral, Dean Eckles, Juanjuan Zhang

Related Coursework in Computer Science, Economics and Management:

6.268 Network Science and Models (Patrick Jaillet)

6.867 Machine Learning (Guy Bresler, Tamara Broderick, Leslie P Kaelbling)

6.S191: Introduction to Deep Learning (Alexander Amini, Ava Soleimany)

14.122 Game Theory (Glenn Ellison)

14.124 Contract Theory (Bengt Holmstrom)

14.137 Behavioral Economics (Drazen Prelec)

14.271 Industrial Organization (Glenn Ellison)

14.273 Structural Industrial Organization (Sarah Ellison, Whitney Newey)

14.282 Organizational Economics (Robert Gibbons, Michael Whinston, John Van Reenen)

14.382 Econometrics (Victor Chernozhukov)

14.387 Applied Econometrics (Joshua Angrist, Victor Chernozhukov)

15.579 Research in Networks (Sinan Aral)

15.840 Experimental Design (Dean Eckles)

New York University (New York, NY)
M.A. Statistics, 2015

University of International Business and Economics (Beijing, China)
B.A. Management, B.A. Economics, 2012

Research

General Interests

Primary: Machine Learning and Field Experiment, Social and Economic Networks.

Secondary: Psychology and Economics.

Selected Work in Progress

Optimizing Targeting Policies via Adaptive Experimentation: Improving User Retention at the Boston Globe (with Dean Eckles, Paramveer Dhillon and Sinan Aral)

Conference Presentations:
INFORMS Marketing Science (Rome, June 2019)
Advances on Field Experiments (University of Chicago, July 2019)
CODE (MIT, November 2019)
NeurIPS CausalML Workshop (Vancouver, December 2019)

We combine offline policy evaluation and optimization with a sequence of contextual bandit experiments to design targeting policy in a proactive churn management setting for the Boston Globe. In the first experiment, we target users with higher churn risk with higher probabilities using a price discount. Based on realized churns from the first experiment, we use Hajek estimator to compute counterfactual outcomes under alternative targeting policies and evaluate their performance. Specifically, we estimate the monetary value of different churn prediction algorithms and subscriber features. We then compute an optimal targeting policy based on estimates of conditional average treatment effects or a doubly robust score. We implement the policy using Thompson Sampling in the second experiment and iterate on it. Our proposed method balances exploration and exploitation and can adapt to potential non-stationarity in the environment.

Does Video Content Predict Consumer Engagement and Product Sales? Evidence from Influencer Advertising on TikTok (with Juanjuan Zhang and Yuhan Zhang)

We use the content of video advertisement on TikTok, a popular social media platform for short video clips, to predict user engagement (measured by the number of likes, shares and comments) and product sales. We first use CNN and LSTM to extract features from the video, then we exploit the fact that these advertising videos are posted at different time to set up a Difference-in-Difference (DID) strategy to estimate the causal effect of advertising videos on sales. Finally we relate video features to the casual effect to investigate what type of content is more effective at generating sales.

The Effect of Verification on Rumor Diffusion

I investigate how uncertainty affects the pattern of information diffusion. In the context of rumors spreading on a social network, I first use a simple model of Bayesian social learning to show that under certain conditions agents have higher incentive to share rumors than non-rumors. This is because when the rumor or public information is noisy, agents' private information is relatively more important, it encourages the dissemination of private information through sharing. And sharing can both be strategic complement and substitute depending on the local network structure. I then analyze a large dataset of over 2 millions tweets that contains over 2400 rumor threads about 5 breaking news on Twitter to validate the model predictions. Exploiting the fact that different rumors got verified at different time points, I use a Difference-in-Difference (DID) approach to compare the diffusion patterns pre and post to verification between rumors that have similar content as measured by word embeddings (Word2vec). I find that verification has a strong and immediate positive effect on number of retweets, and the effect is more pronounced for true rumors. The effect attenuates over time and eventually becomes significantly negative, that is, verification reduces the number of retweets in the future comparing with rumors that are not yet verified.

Research Assistantship

Massachusetts Institute of Technology

Prof. Sinal Aral: Field Experiments on Connected Stores and Smart Retailing (2015-2016)

New York University

Prof. Vishal Singh: Mining and Visualizing Large Text Corpus (2014-2015)

Prof. Jennifer Hill: Bias Amplification in Hierarchical Models (2014-2015)
Middleton et al. (2016) "Bias Amplification and Bias Unmasking" (*Political Analysis*)

Prof. Patrick Sharkey:
Field Experiments on the Effect of Neighborhood Stigma on Economic Transaction (2013-2014)
Besbris et al. (2015) "The Effect of Neighborhood Stigma on Economic Transactions" (*PNAS*)

Teaching Assistantship

Massachusetts Institute of Technology

Prof. Sinal Aral: Digital Marketing (MBA) Fall 2017

Prof. Roberto Rigobon: Macroeconomic Policy (MBA) Summer 2018

Prof. David Gamarnik: Data, Model and Decision (MBA) Summer 2018, 2019

Prof. John Little: Singularity (MBA) Fall 2018

Prof. Birger Wernerfelt: Marketing Strategy (MBA) Fall 2018

Prof. Dean Eckles: Experimental Design and Analysis (PhD) Fall 2019

Prof. Erik Brynjolfsson: Analytics Lab (MBA) Summer 2019

Language and Computational Skills

Languages: Mandarin (native), English (fluent).

Programming: R, Python/Keras.

Summary of Interests in Economics of Digitization

My interest in the economics of digitization lies in four broad areas:
(1) applying machine learning techniques to facilitate casual inference via the identification of heterogeneous treatment effects (e.g., casual forest) to learn an optimal personalized decision/treatment rules.
(2) machine learning aided analysis of unstructured data such as text, image and video.
(3) the economic value of consumer data, for example, what's the value of a targeting policy using certain consumer features vs. not using them.
(4) how consumer judgment and decision making can be affected by AI systems, and particularly the interplay between algorithmic interpretability and trust.