Research Statement

Bruno Abrahao
New York University
bd58@nyu.edu

1 Overview

My research program lies at the intersection of computational social science, network analysis, and machine learning. Across two decades, I have investigated how the structural and algorithmic properties of networked systems shape human behavior, and how computational methods can be designed to operate reliably in complex social and technological environments. This agenda spans six interconnected areas: (1) trust, reputation, and digital platforms; (2) urban mobility, segregation, and public health; (3) machine learning robustness and deployment; (4) computational social science, digital health, and organizations; (5) network inference and community structure; and (6) Internet measurement and web performance. My work has appeared in leading outlets including *Proceedings of the National Academy of Sciences* (*PNAS*), *KDD*, *ICML*, *CHI*, *AAAI*, *Network Science*, *Sociological Science*, and *Socius*, among others.

2 Trust, Reputation, and Digital Platforms

A central and long-running thread of my research investigates how reputation systems mediate trust in online marketplaces, with sustained attention to the sharing economy. The cornerstone of this program is a large-scale field experiment on Airbnb, conducted with Parigi, Gupta, and Cook (Abrahao et al., 2017). We demonstrated that platform reputation can offset the influence of social biases: guests with strong reviews were equally likely to be accepted regardless of their race, whereas guests without reputational signals experienced significant discrimination. Published in *PNAS*, this study provided some of the earliest causal evidence that algorithmic reputation mechanisms can function as bias-correction devices in two-sided markets.

This experiment was informed by earlier theoretical and empirical groundwork. State, Cook, and I first examined how power imbalances in online exchange shape cooperative behavior (State et al., 2012), work that received the Best Extended Abstract Award at WebSci 2012. We subsequently showed that such asymmetries systematically distort rating behavior, producing structural biases in reputation scores (State et al., 2016). Qiu, Parigi, and I then examined the differential effects of reputation signals—star ratings versus review count—on trust formation, finding that the marginal effect of an additional star diminishes faster than that of an additional review, with direct implications for platform interface design (Qiu et al., 2018).

Lu and I extended this investigation to the structural and temporal dimensions of rating bias, analyzing how biases such as homophily and power dependence co-evolve with reputational systems at the population level (Lu and Abrahao, 2019; Abrahao and Lu, 2019). In parallel, Zhu, Zhao, Parigi, and I drew causal conclusions from a longitudinal experiment tracking Airbnb users over time, finding that positive interactions enhance platform trust while reducing the weight users place on social similarity (Zhu et al., 2020).

Two ongoing projects extend this line. Huang, Parigi, Zhu, and I are studying how personal transaction experiences reinforce or update trust judgments over repeated interactions on Airbnb (Abrahao et al., 2026a). Halaburda, Chai, and I examine the erosion of social capital in digital exchange platforms by analyzing the structural decline of Couchsurfing (Abrahao et al., 2026b, 2025a).

3 Urban Mobility, Segregation, and Public Health

A more recent line of work, in collaboration with Marlow and Makovi, applies network methods to large-scale mobility data to understand urban segregation and its consequences for public health. Using smartphone-derived mobility flows, we documented the sharp increase in neighborhood isolation during the early months of the COVID-19 pandemic and showed that isolation varied systematically with neighborhood socioeconomic composition (Marlow et al., 2021). Published in *Sociological Science*, this was among the first studies to quantify pandemic-induced changes in inter-neighborhood connectivity at national scale.

A follow-up study published in *Socius* demonstrated that these changes were remarkably durable: as of 2022, many U.S. cities had not recovered pre-pandemic levels of cross-neighborhood movement (Marlow et al., 2023). György, Marlow, Makovi, and I further showed that pre-existing patterns of segregated mobility amplified neighborhood-level disparities in COVID-19 spread (György et al., 2023). Published in *Network Science*, this study established that mobility network structure is a first-order determinant of both social cohesion and disease transmission in cities.

4 Machine Learning Robustness and Deployment

My machine learning research focuses on improving the reliability of classifiers deployed in environments that differ from their training distribution. Wang, Kamar, and I developed a supervised method for discovering *unknown unknowns*—coherent input regions where a deployed classifier fails silently—by mining test samples that are systematically misclassified (Wang et al., 2020). Published at AAAI 2020, this approach identifies failure modes without requiring prior knowledge of what the model might get wrong. Wang, Ahmed, Zhu, and I extended this into a model rectification framework that extracts unknown-unknown clusters from deployment data and uses them to fine-tune the original model, reducing error rates under distribution shift without sacrificing in-distribution accuracy (Abrahao et al., 2021a).

Two current projects advance this agenda. Wang, Qin, Valle, and I are developing BootOOD, a self-supervised approach to out-of-distribution detection that leverages synthetic sample exposure under the neural collapse phenomenon to improve detection without requiring real OOD training examples (Wang et al., 2026). Zhou, He, Wang, Chen, and I have addressed the complementary problem of data effectiveness assessment for user churn detection, proposing an indicator-assisted framework that unifies multiple data quality dimensions (Zhou et al., 2026). This work, forthcoming in *Tsinghua Science and Technology*, bridges machine learning methodology and business analytics.

In a related direction, Bello, Bobrov, Li, Traina, and I are exploring the use of large language models for measuring and valuing service quality at scale (Abrahao et al., 2026c). By applying LLMs to extract structured quality assessments from unstructured customer interactions, we aim to enable real-time service monitoring that captures dimensions—such as empathy, accuracy, and responsiveness—that traditional metrics miss.

5 Computational Social Science, Digital Health, and Organizations

A cross-cutting stream of my research leverages digital trace data to study social and health phenomena. Saha, Sugar, Torous, Kiciman, De Choudhury, and I used social media data to study the behavioral effects of psychiatric medication use, finding measurable changes in language, activity, and social engagement patterns following self-reported medication initiation (Saha et al., 2019). This study received the Outstanding Study Design Award at ICWSM 2019 and demonstrated the potential of passive digital phenotyping as a complement to clinical observation.

Parigi and I synthesized methodological perspectives across these lines in a chapter for the *Oxford Handbook of Social Networks*, reviewing how computational social science, big data, and network methods are reshaping the study of social structure (Abrahao and Parigi, 2021). The chapter examines both the opportunities afforded by large-scale digital data and the methodological challenges—including selection bias, measurement validity, and ethical considerations—that researchers must navigate.

Li, Rider, and I have been investigating the reciprocal relationship between cultural and structural embeddedness in organizations (Abrahao et al., 2026d, 2025b). Using natural language processing to measure cultural alignment from internal communications, we link cultural fit to network positioning and career outcomes. This work is currently under review at *Management Science*.

6 Network Inference and Community Structure

A substantial body of my work addresses the algorithmic foundations of reasoning about network structure from incomplete data. Soundarajan, Hopcroft, Kleinberg, and I introduced a *separability framework* for community detection, providing theoretical conditions under which community structure can be reliably distinguished from random fluctuations (Abrahao et al., 2012). The initial conference paper at KDD 2012 was later expanded into a full treatment in *ACM Transactions on Knowledge Discovery from Data* (Abrahao et al., 2014), unifying several existing quality measures and providing a principled basis for comparing community detection algorithms.

Complementing this structural perspective, Chierichetti, Kleinberg, Panconesi, and I studied the *trace complexity* of network inference—the minimum number of cascade observations needed to reconstruct an underlying network (Abrahao et al., 2013). Published at KDD 2013, this work established information-theoretic lower bounds and showed that recovery is fundamentally harder when cascades are short or the network is dense, with implications for epidemiological and information-diffusion studies.

During this period, I also contributed to online optimization in network settings. Lin, Kleinberg, Lui, Chen, and I formulated a combinatorial partial monitoring game with linear feedback and demonstrated its application to problems such as influence maximization under uncertainty (Lin et al., 2014). Published at ICML 2014, this work extended the partial monitoring framework to combinatorial action spaces, providing improved regret bounds.

7 Internet Measurement and Web Performance

My earliest research examined the statistical structure of Internet traffic and web workloads. In collaboration with Menascé, Barbará, Almeida, and Ribeiro, I showed that web request streams exhibit fractal (self-similar) properties with direct implications for capacity planning (Menascé et al., 2002a,b). This work, a Best Paper Award finalist at The Web Conference (WWW 2002), established that heavy-tailed distributions in session behavior render traditional Poisson-based models inadequate. In a related study, Fonseca, Almeida, Crovella, and I demonstrated that temporal correlations in web access patterns could be exploited for caching and prefetching (Fonseca et al., 2003).

These insights motivated practical contributions to resource management. Zhang and I developed methods for characterizing application workloads on CPU utilization for utility computing environments (Abrahao and Zhang, 2004), work that led to a U.S. patent on computational analysis methods (Abrahao and Zhang, 2010). With Almeida, Zhang, Beyer, and Safai, I developed a self-adaptive framework for SLA-driven capacity management in Internet services (Abrahao et al., 2006). Moving from the application layer to network-layer measurement, Kleinberg and I studied the dimensionality of the Internet delay space, showing that the metric structure of round-trip times admits low-dimensional embeddings useful for network coordinate systems (Abrahao and Kleinberg, 2008a,b). This work appeared at both the ACM Internet Measurement Conference and the ACM Symposium on Principles of Distributed Computing in 2008.

8 Future Directions

Looking ahead, I plan to deepen the connections across these research streams, particularly at the intersection of platform design, algorithmic fairness, and organizational networks. The increasing embeddedness of AI in economic and social systems creates both new measurement opportunities and new risks, and I believe the combination of network methods, causal inference, and machine learning that has characterized my work is well positioned to address these challenges.

References

Abrahao, B., Parigi, P., Gupta, A., and Cook, K. (2017). Reputation offsets trust judgments based on social biases among Airbnb users. *Proceedings of the National Academy of Sciences*, 114(37):9848–9853. doi:10.1073/pnas.1604234114.

State, B., Abrahao, B., and Cook, K. (2012). From power to status in online exchange. In *Proceedings of the 4th ACM International Conference on Web Science (WebSci)*, pages 291–300, Evanston, IL. ACM.

State, B., Abrahao, B., and Cook, K. (2016). Power imbalance and rating systems. In *Proceedings of the 10th International AAAI Conference on Web and Social Media (ICWSM)*, pages 368–377, Cologne, Germany. AAAI Press. doi:10.1609/icwsm.v10i1.14753.

Qiu, W., Parigi, P., and Abrahao, B. (2018). More stars or more reviews? Differential effects of reputation on trust in the sharing economy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montreal, Canada. ACM. doi:10.1145/3173574.3173727.

Lu, A. and Abrahao, B. (2019). The structure and evolution of online rating biases in the sharing economy. In *Proceedings of the Pacific Asia Conference on Information Systems (PACIS)*, Xi'an, China.

Abrahao, B. and Lu, A. (2019). The structure and evolution of online rating biases in the sharing economy. In *15th Annual Symposium on Statistical Challenges in Electronic Commerce Research (SCECR)*, Hong Kong.

Zhu, C., Abrahao, B., Zhao, A., and Parigi, P. (2020). The effect of user interactions on shaping online trust: Evidence from a large-scale experiment. In *Proceedings of the Pacific Asia Conference on Information Systems (PACIS)*, Dubai, UAE.

Abrahao, B., Huang, S., Parigi, P., and Zhu, C. (2026). Personal experiences reinforce trust judgments: A field study on Airbnb. Working paper, New York University.

Abrahao, B., Halaburda, H., and Chai, H. S. (2026). The erosion of social capital in digital exchange: Evidence from the decline of Couchsurfing. Working paper, New York University.

Abrahao, B., Halaburda, H., and Chai, H. S. (2025). The erosion of social capital in exchange platforms: How socio-psychological biases shaped Couchsurfing's downfall. In *Statistical Challenges in Electronic Commerce Research (SCECR)*, Singapore.

Marlow, T., Makovi, K., and Abrahao, B. (2021). Neighborhood isolation during the COVID-19 pandemic. *Sociological Science*, 8(9):170–190. doi:10.15195/v8.a9.

Marlow, T., Makovi, K., and Abrahao, B. (2023). Durable change in U.S. urban mobility networks, 2019–2022. *Socius*, 9. doi:10.1177/23780231231198857.

György, A., Marlow, T., Abrahao, B., and Makovi, K. (2023). Segregated mobility patterns amplify neighborhood disparities in the spread of COVID-19. *Network Science*, 11(3):411–430. doi:10.1017/nws.2023.6.

Wang, Z., Abrahao, B., and Kamar, E. (2020). Supervised discovery of unknown unknowns through test sample mining. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, pages 13959–13960, New York, NY. AAAI Press. doi:10.1609/aaai.v34i10.7252.

Abrahao, B., Wang, Z., Ahmed, H., and Zhu, Y. (2021). Model rectification via unknown unknowns extraction from deployment samples. *arXiv preprint arXiv:2102.04145*.

Wang, Y., Qin, T., Valle, E., and Abrahao, B. (2026). BootOOD: Self-supervised out-of-distribution detection via synthetic sample exposure under neural collapse. Working paper.

Zhou, M., He, Q., Abrahao, B., Wang, X., and Chen, Y. (2026). Unifying data effectiveness assessment in user churn detection: An indicator-assisted framework. *Tsinghua Science and Technology*. doi:10.26599/TST.2024.9010230.

Abrahao, B., Bello, O., Bobrov, A., Li, X., and Traina, J. (2026). Measuring and valuing service quality with large language models. Working paper, New York University.

Saha, K., Sugar, B., Torous, J., Abrahao, B., Kiciman, E., and De Choudhury, M. (2019). A social media study on the effects of psychiatric medication use. In *Proceedings of the 13th International AAAI Conference on Web and Social Media (ICWSM)*, pages 440–451, Munich, Germany. AAAI Press. doi:10.1609/icwsm.v13i01.3242.

Abrahao, B. and Parigi, P. (2021). Computational social science, big data, and networks. In Light, R. and Moody, J., editors, *The Oxford Handbook of Social Networks*, pages 516–534. Oxford University Press. doi:10.1093/oxfordhb/9780190251765.013.23.

Abrahao, B., Li, X., and Rider, C. (2026). Fitting out and standing in? The reciprocal nature of cultural and structural embeddedness in organizations. Working paper, New York University.

Abrahao, B., Li, X., and Rider, C. (2025). Enculturation and intermediation: Mechanisms of career attainment within organizations. In *People and Organizations Conference*, Wharton AI & Analytics Initiative, University of Pennsylvania.

Abrahao, B., Soundarajan, S., Hopcroft, J., and Kleinberg, R. (2012). On the separability of structural classes of communities. In *Proceedings of the 18th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 624–632, Beijing, China. ACM. doi:10.1145/2339530.2339631.

Abrahao, B., Soundarajan, S., Hopcroft, J., and Kleinberg, R. (2014). A separability framework for analyzing community structure. *ACM Transactions on Knowledge Discovery from Data*, 8(1):1–30. doi:10.1145/2527231.

Abrahao, B., Chierichetti, F., Kleinberg, R., and Panconesi, A. (2013). Trace complexity of network inference. In *Proceedings of the 19th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 491–499, Chicago, IL. ACM. doi:10.1145/2487575.2487664.

Lin, T., Abrahao, B., Kleinberg, R., Lui, J., and Chen, W. (2014). Combinatorial partial monitoring game with linear feedback and its applications. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 901–909, Beijing, China. PMLR.

Menascé, D. A., Abrahao, B., Barbará, D., Almeida, V., and Ribeiro, F. (2002a). Fractal characterization of web workloads. In *Proceedings of the 11th International Conference on the World Wide Web (WWW)*, Honolulu, HI.

Menascé, D. A., Abrahao, B., Barbará, D., Almeida, V., and Ribeiro, F. (2002b). Characterizing e-business workloads using fractal methods. *Journal of Web Engineering*, 1(1):74–90.

Fonseca, R., Almeida, V., Crovella, M., and Abrahao, B. (2003). On the intrinsic locality properties of web reference streams. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 448–458. IEEE. doi:10.1109/INFCOM.2003.1208696.

Abrahao, B. and Zhang, A. (2004). Characterizing application workloads on CPU utilization for utility computing. *HP Labs Technical Reports*, HPL-2004-157, Palo Alto, CA.

Abrahao, B., Almeida, V., Almeida, J., Zhang, A., Beyer, D., and Safai, F. (2006). Self-adaptive SLA-driven capacity management for Internet services. In *Proceedings of the 10th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, pages 557–568. IEEE. doi:10.1109/NOMS.2006.1687584.

Abrahao, B. and Kleinberg, R. (2008a). On the Internet delay space dimensionality. In *Proceedings of the 8th ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 157–168, Vouliagmeni, Greece. ACM. doi:10.1145/1452520.1452541.

Abrahao, B. and Kleinberg, R. (2008b). On the Internet delay space dimensionality. In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 259–268. ACM. doi:10.1145/1400751.1400810.

Abrahao, B. and Zhang, A. (2010). Method and apparatus for computational analysis. U.S. Patent #7,698,111.