SIGecom Winter Meeting 2025 Highlights

BAHAR BOROOMAND
University of Alberta
and
SAFWAN HOSSAIN
Harvard University
and
EDEN SAIG
Technion – Israel Institute of Technology

Bahar Boroomand is a M.Sc. student in Computing Science at the University of Alberta, passionate about Machine Learning and its applications. Her current research focuses on alleviating biases caused by rating-based scoring algorithms in recommender systems using data-driven machine learning techniques.

Safwan Hossain is a PhD Candidate in Computer Science at Harvard University. His research interests are broadly at the intersection of economics and computer science, involving questions related to strategic behavior, fairness, and incentives that arise in supervised or online learning settings with multiple agents. Prior to his PhD, Safwan received his BASc. in Electrical Engineering and MSc. in Computer Science from the University of Toronto, and spent two years working as a machine learning engineer at Cerebras Systems.

Eden Saig is a PhD candidate in Computer Science at the Technion, advised by Nir Rosenfeld. His research focuses on machine learning and algorithmic decision-making in social contexts, aiming to develop socially favorable learning algorithms for behavioral environments with dynamics and incentives. Before starting his PhD, Eden received a BSc in Computer Science, BSc in Physics, and an MSc in Computer Science, all from the Technion.

General Terms:

Additional Key Words and Phrases:

---

1. INTRODUCTION TALKS

1.1 Haifeng Xu: Rethinking Online Content Ecosystems through the Lens of Computational Economics

The first invited talk of the session, by Haifeng Xu from the University of Chicago, highlighted a new research agenda: studying the wide range of problems in online content ecosystems through the formalisms of computational economics. Online content recommendation engines—core to platforms like YouTube, Instagram, and TikTok—serve personalized content to billions of users daily. The classic model considers both the users and the content library to be static, with the recommendation engine responsible for generating a mapping between the two. Xu's talk envisions a richer model that incorporates the incentives of content creators (e.g., YouTube rewarding videos based on length and views), the myopic and dynamic behavior of consumers, and the increasingly prominent role of AI in both generating content and being trained on it. This is a rich, dynamic multi-agent environment and the remainder of the talk considers two distinct directions within this framework:

(1) Diagnosing and optimizing existing content ecosystems

(2) How AI-generated content can transform future content ecosystems

Existing content ecosystems can be seen as a two-sided market between self-interested consumers and creators, with the platforms acting as a powerful and self-interested intermediary. Recent works have studied parts of this interaction. Several recent works study the "supply-side" interactions between creators, who generate traffic, and platforms, who benefit from this traffic and share revenue. These include understanding, among others, creator competition [Ben-Porat and Tennenholtz 2018; Ben-Porat et al. 2020], incentivized matching mechanisms [Mladenov et al. 2020], and content distribution at equilibrium [Jagadeesan et al. 2023]. Less studied is the "demand-side", which model interactions between platforms and users. [Kleinberg et al. 2024] focus on improving recommendations through "behavior-aware" system learning. Modeling the overall ecosystem with all three types of players is understudied. Of note is [Yao et al. 2023]: they study mechanisms to incentivize content creation for user welfare maximization under a self-interested platform. The proposed mechanism ends up introducing more competition for congested topics. Importantly, a variant of the mechanism was tested and validated on Instagram Reels, involving over ten million users and creators.

Xu suggests that the rise of powerful AI systems constitutes a *fourth* player within this ecosystem. AI systems can act as content generators and thus compete with human creators. In turn, they also rely on platforms and user feedback to train and validate their models. Each of these roles/perspectives lead to numerous unexplored research questions and can fundamentally alter the dynamics of content ecosystems. [Taitler and Ben-Porat 2025], for instance, study the AI-creator-platform dynamic and suggests that AI systems can strategically give worse answers to allow for more high-quality human generated content in the short term in order to increase their long-term utility. They also observe a Braess paradox phenomenon occurring once AI systems partake in content generation. [Raghavan 2024] studies the AI-consumer interaction and suggests that it may lead to reduced content diversity. [Duetting et al. 2024] studies the AI-platform-consumer interaction and illustrates how AI systems can be part of new monetization mechanisms. Overall, the talk concludes by stressing that "incentives and agency are crucial to both learning algorithms and market mechanisms for resolving these pressing issues".

1.2 Jon Kleinberg: Language Generation in the Limit

In his intriguing talk, Jon Kleinberg presented a formal abstraction which aims to capture the foundational properties of generative AI [Kleinberg and Mullainathan 2024]. He began by asking whether there exists a simple theoretical metaphor — analogous to the metaphor of "Alice and Bob" in secure communication, or the metaphor of "Byzantine generals" in distributed systems — which captures the core properties of generative AI and enables rigorous analysis. Towards this, Kleinberg proposed framing the task of "learning to generate" as an algorithmic question:

Can an algorithm, presented only with a stream of valid words of some formal language, eventually start emitting never-before-seen words of that language?

To formalize this question, [Kleinberg and Mullainathan 2024] extend the classical framework of *language learning in the limit*, initially formulated by [Gold 1967], and further characterized by [Angluin 1979; 1980]. In the framwork, a language L is a countably infinite set of words, and there is a countable set of languages $\{L_1, L_2, \dots\}$. An adversary initially selects a target language L_k from that set, and interacts with an algorithm over discrete time steps. At each step, the adversary reveals a previously unseen word $w \in L_k$, the algorithm emits an output, and no further information about L_k is provided except for these positive examples.

In the original Gold-Angluin framework, the output of the algorithm at each step is a guess about the index of the target language L_k, and the goal is design an algorithm which stops making mistakes after a finite number of steps. The classic result of [Gold 1967] shows that this task is impossible in general, as an adversary could construct word streams for which the algorithm makes an infinite number of mistakes. However, when shifting focus from language identification to language generation, [Kleinberg and Mullainathan 2024] reveal a fundamental contrast: They present an algorithm that, in the limit, produces an infinite stream of valid and previously unseen strings from the target language L_k, despite not being able to explicitly identify it in the Gold-Angluin sense.

The algorithm relies on the definition of *language criticality*, which identifies progressively thinner languages consistent with the data seen so far. At each step, the generation algorithm maintains the critical language, and generates a previously-unseen word from it. While this guarantees validity in the limit, the definition of criticality also implies that each critical language is a strict subset of the previous ones. Thus, the algorithm may reach a state where the critical language is a strict subset of L_k, preventing it from generating all possible words in the target language. This reveals a trade-off between validity and breadth: to avoid mistakes in generation, the algorithm must permit incomplete coverage of the target language. Interestingly, this trade-off draws qualitative parallels to linguistic phenomena observed in practice, such as vernacular adoption dynamics in online communities [Danescu-Niculescu-Mizil et al. 2013], and quality-diversity tradeoffs in LLMs.

Beyond their main result, [Kleinberg and Mullainathan 2024] provide stronger convergence guarantees for finite sets of languages, and extend the framework to settings with prompting. Subsequent work has already begun exploring different aspects of the validity–breadth trade-off [Charikar and Pabbaraju 2024; Kalavasis et al. 2024a; 2024b; Kleinberg and Wei 2025], extending the stronger convergence guarantees to certain infinite sets of languages [Li et al. 2024], and exploring interaction models with noisy examples [Raman and Raman 2025]. Each line of inquiry provides new perspectives on the fundamental properties of language generation, and creates intriguing frontiers for future work.

1.3 Manish Raghavan: Competition and Diversity in Generative AI

In his talk, Manish Raghavan explored the tension between competition and diversity in the context of generative AI, drawing attention to a growing concern: generative models become ubiquitous across many domains, but the outputs they produce remain relatively homogeneous. This phenomenon, which relates the the general notion of *algorithmic monoculture* [Kleinberg and Raghavan 2021], arises when many individuals rely on the same language model, leading to results which are less diversified. For instance, while AI tools may enhance individual productivity in tasks such as brainstorming or ideation, they might also increase homogeneity by guiding users toward similar answers. This motivates a natural question: Can we design environments that encourage novelty alongside correctness?

Towards this, [Raghavan 2024] introduces a stylized game-theoretic model to study this question. The model defines a game over n players, where each action is a categorical distribution over outputs with an ordering constraint, representing the output distribution of an LLM given some prompt. When multiple players have the same realized output, they split the reward, reflecting competition for audience attention or market share. The theoretical analysis shows that stronger competition induces players to adopt more diverse strategies, although equilibrium behavior remains less diverse than the social optimum. Perhaps surprisingly, the relative ranking of different strategies depends on competitive intensity, and a generative model that has the best performance in isolation can become suboptimal in the presence of competition due to lack of diversity.

Empirical validation is performed through simulations of the game Scattergories, played by LLMs under two settings: one where players share the same language model but choose generation temperatures strategically, and another where they can also choose which model to use. The results demonstrate that the best sampling strategy depends not only on the temperature but also on the specific model and the number of players. Models better at sampling from the tails of their output distributions had greater diversity in their outputs, and performed better as competitive pressure increased.

The talk concluded with several takeaways and open questions. While generative AI tools hold immense promise, their widespread adoption risks diminishing diversity. Competition, both between users and between models, can act as a force to counteract this. This points to a broader design question for AI systems: We often optimize systems for correctness, but can we optimize for novelty and diversity? As AI-generated content permeates more aspects of society, understanding and shaping these dynamics will be a vital challenge for both theorists and practitioners.

1.4 Yannai Gonczarowski: Algorithmic Collusion by Large Language Models

Yannai Gonczarowski presented a talk on his recent paper of the same title, coauthored with Sara Fish and Ran Shorrer [Fish et al. 2024]. He begins by defining the classical notion of collusion in economics: traders/competitors meeting to jointly raise the price of a certain good, at the expense of the public. He comments that in an increasing number of settings, automated AI driven agents are being used for pricing, and this work formally explores the potential for autonomous algorithmic collusion when large language models (LLMs) are used for this task. Three main questions are addressed within this context:

(1) Are LLMs good at pricing tasks?
(2) If multiple firms separately use LLMs for pricing, can this lead to supracompetitive prices?
(3) What mechanisms promote or prevent collusion?

Using a repeated Bertrand oligopoly environment, the authors first demonstrate that LLMs are capable of learning near-optimal monopoly pricing quickly and reliably in a monopoly setting. The pricing agent here does not require any fine-tuning and is instead based completely on in-context information. This includes the task prompt, basic information about the instance, the market history (past prices, quantities sold, profits earned etc), and past reasoning stated by the LLM. GPT-4 converges to the optimal monopoly price in all settings, with other LLMs showing more varied performances.

The duopoly setting is considered next with the precise question: if two firms are using GPT-4 for pricing, does it lead to competitive or supra-competitive pricing? Naturally, the in-context prompt used for pricing matters deeply, and the work considers two variants: (1) explicitly mentioning the LLM to not undermine profitability, and (2) mentioning that lower pricing than competitors will lead to higher sales volume. While both prompts lead to supra-competitive pricing, prompt (1) leads to higher prices than prompt (2). To understand the pricing process further, the LLM reasoning is analyzed, and evidence illustrates that the agents are concerned about avoiding a price war, especially under prompt (1).

The talk underscores the regulatory challenges posed by LLMs: their ability to autonomously adopt collusive strategies even under benign instructions, their black-box reasoning, and the sensitivity of outcomes to prompt wording. Unlike traditional Q-learning agents, the basis of past automated pricing works, LLMs are pre-trained, adaptable, and readily deployable, exacerbating concerns over the real-world applicability of algorithmic collusion. Reevaluating regulatory frameworks in light of these findings is fundamental.

1.5 Sanmi Koyejo: On Shaping AI's Impact on Billions of Lives

In his talk, Sanmi Koyejo presents a multifaceted vision for developing artificial intelligence technologies that maximize societal benefit while mitigating harm. He argues for reorienting AI development toward the public good by embedding economic, ethical, and sociotechnical considerations into the design and deployment of AI systems. The underlying premise is that the default trajectory of AI, driven largely by market incentives, may not align with broader societal interests unless interventions are made deliberately and early.

A core theme of the talk is the human-AI collaboration paradigm. Rather than envisioning AI as a replacement for human labour, Koyejo advocates for building synergistic systems that augment human capabilities, improve job satisfaction, and unlock elastic economic potential. For example, in domains like consulting, legal services, and writing, early empirical evidence suggests that AI disproportionately benefits lower-skilled professionals by narrowing performance gaps. Importantly, the speaker draws on economic theory to emphasize that AI's impact on employment will vary by sector, depending on how demand responds to increased efficiency. In areas where greater efficiency leads to increased usage, such as healthcare or education, AI has the potential to create more jobs by expanding services. In contrast, in sectors where demand remains relatively fixed, such as agriculture, efficiency gains are more likely to reduce the number of workers needed.

The talk highlights several concrete application domains where AI can be transformative. In healthcare, he discusses the potential for AI to alleviate administrative drudgery and reduce burnout among clinicians, enabling them to focus more on patient care. In education, AI can help close systemic learning gaps through personalized instruction and empirically driven interventions. Koyejo stresses the importance of continuous evaluation and measurement infrastructure in both domains, drawing parallels to high-stakes fields like clinical trials. A key insight is the need for a shift from static, pre-deployment evaluation to dynamic, post-deployment monitoring, a critical requirement given the evolving nature of AI systems and their societal impacts.

He also devotes considerable attention to the information ecosystem, where trust, polarization, and misinformation present pressing challenges. He identifies the dangers of overtrust in AI-generated content, especially in the context of natural language interfaces, and proposes mechanisms for calibrated trust, such as user-facing confidence indicators, citation linking, and interpretable model diagnostics.

Rather than prescribing a single moonshot, Koyejo calls for a portfolio of milestone-driven efforts, ranging from targeted prize challenges to the establishment of interdisciplinary research centers. He encourages researchers to contribute to foundational infrastructure. This pluralistic approach reflects the belief that shaping AI's impact on billions requires collective, iterative innovation rather than top-down mandates. Policymakers, technologists, and civil society actors are urged to co-create governance mechanisms that are legally grounded yet adaptable to the unique demands of AI.

2. FIRESIDE CHAT WITH PRESTON MCAFEE AND PRABHAKAR RAGHAVAN

The 2025 Winter Meeting featured a thought-provoking Fireside Chat between Preston McAfee, Google Distinguished Scientist and a pioneering expert in auctions, market design, and computational economics, and Prabhakar Raghavan, Google's Chief Technologist and a renowned authority on search, algorithms, and web-scale systems. Drawing on decades of influential research and leadership across academia and industry, McAfee and Raghavan engaged in a dynamic conversation about how AI could influence market behavior, support proof automation in microeconomic theory, expose limitations in current macroeconomic modelling, and introduce new approaches to reasoning about complex socio-economic systems.

Preston McAfee, Google Distinguished Scientist, is an expert on pricing, auctions, antitrust, business strategy, market design, computational advertising, and machine learning applied to exchanges. He has published over 130 refereed articles, holds eleven patents, and has authored three books. His research notably influenced spectrum auction design, earning him the Golden Goose award. After earning his B.A. from the University of Florida and his Ph.D. in economics from Purdue University, McAfee spent 28 years as a professor at UWO, UT Austin, and Caltech. He held leadership roles at Yahoo!, Google, and Microsoft, including Chief Economist at Microsoft. In 2006, he published the open-access textbook Introduction to Economic Analysis, awarded the SPARC Innovator Award in 2009.

Prabhakar Raghavan is Google's Chief Technologist and one of the foremost authorities on algorithms and web search. He is the co-author of the foundational texts Randomized Algorithms and Introduction to Information Retrieval. Prabhakar has published over 100 papers and holds 20 patents, particularly in link analysis. At Google, he served as Senior VP for Knowledge & Information, overseeing products like Search, Ads, and Gemini. Before Google, he led Yahoo! Labs, served as CTO of Verity, and spent 14 years at IBM Research. He holds a Ph.D. from UC Berkeley and a B.Tech from IIT Madras. Prabhakar is a member of the National Academy of Engineering, a Fellow of ACM and IEEE, and recipient of the 2017 WWW Test of Time Award.

How do you think AI will influence markets beyond traditional concerns like collusion?

Prabhakar. One promising area lies in using reinforcement learning to augment mathematical proofs, particularly in theoretical computer science and microeconomics. Recent efforts have explored automating proof discovery for hardness of approximation results by learning to optimize the "gadgets" that underpin such proofs. While these AI systems haven't yet proven new theorems, they have independently rediscovered known ones using novel constructions not present in training data. These methods, if extended to combinatorial auctions and mechanism design, may refine classical hardness results and reduce proof complexity, but they still face challenges in validation and formal proof checking.

Preston. The macroeconomic side presents deeper methodological challenges. Traditional economic counterfactuals, such as Fogel's analysis of GDP without railroads, relied on constructing plausible substitutes to estimate upper bounds. This logic can be adapted to AI's economic impact by asking what it would cost to replicate AI's functionality via non-AI means, but it remains flawed, since the bundle of activities changes when costs drop. AI shifts the equilibrium by enabling behaviors that weren't previously feasible. The difficulty lies in modelling these dynamic substitutions and interdependencies across sectors.

Can LLMs or RL-based systems meaningfully contribute to macroeconomic modeling?

Prabhakar. While LLMs encapsulate vast textual knowledge, they reflect how people write about behavior rather than how they act. This makes them imperfect for modelling human strategy. However, at the aggregate level, macro behavior is often smoothed out, allowing some usefulness in high-level prediction. Drawing inspiration from DeepMind's trajectory, from playing Atari to solving protein folding, Prabhakar suggested that macroeconomic simulations could eventually be framed as multi-agent reinforcement learning environments. Agents could evolve over repeated rounds, discovering stable strategies akin to economic equilibria.

Preston. Indeed, modelling economies requires hybrid systems, treating some actors as markets and others, like corporations or key individuals, as decision-makers. While firms understand their supply chains, they struggle to model systemic interdependencies. Here, AI could help by simulating how individual decisions propagate through complex global trade networks. With trade integration now twice as deep as in 1928, understanding these chains is essential, especially amid rising protectionism. AI may offer the only scalable way to capture such emergent, nonlinear effects.

How might qualitative or textual data enhance traditional economic models?

Prabhakar. Language models can act as transformation systems: converting internal events or raw data into polished narratives, and potentially reversing that process. This opens up opportunities to extract soft signals, like executive churn or tone of announcements, and feed them into economic forecasts. For example, capturing why certain cars under- or over-perform in sales, despite technical specifications, may hinge on media coverage and public perception. These hybrid models blending structured and unstructured data could redefine how economists model demand or investor response.

Preston. Traditional models underweight soft signals because they are hard to quantify. Events like major leadership changes currently impact stock prices through gut reactions. But LLMs offer a path to formalizing these signals. Mapping unstructured news into structured risk assessments or demand adjustments could allow for richer, more sensitive models. This is especially valuable in markets where sentiment and narrative matter as much as measurable fundamentals.

What are the broader implications for modelling and equilibrium analysis in AI-influenced systems?

Preston. In practice, economic behavior often deviates from equilibrium. At Yahoo, advertising markets rarely settled into static outcomes. Instead, they evolved through reactive strategies. This mirrors the behavior of generative adversarial networks, which approximate equilibria not through optimization but via iterative best responses. Evolutionary dynamics, rather than rational-agent assumptions, may offer more realistic economic models, albeit harder to construct. These models could better fit real data and support policy decisions in complex, adaptive systems.

Prabhakar. Looking ahead, the key question is whether AI-infused macroeconomic analysis will influence actual policymaking. It's one thing to critique policy papers using LLMs or propose speculative models; it's another to change how governments approach economic strategy. The hope is that in the next 10–15 years, these tools won't just enrich analysis but reshape how economic decisions are made, grounding policy in richer, AI-assisted modelling that bridges qualitative and quantitative domains.

REFERENCES

Angluin, D. 1979. Finding patterns common to a set of strings. In *Proceedings of the eleventh annual ACM Symposium on Theory of Computing*. 130–141.

Angluin, D. 1980. Inductive inference of formal languages from positive data. *Information and control 45,* 2, 117–135.

Ben-Porat, O., Rosenberg, I., and Tennenholtz, M. 2020. Content provider dynamics and coordination in recommendation ecosystems. *Advances in Neural Information Processing Systems 33*, 18931–18941.

Ben-Porat, O. and Tennenholtz, M. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems 31*.

Charikar, M. and Pabbaraju, C. 2024. Exploring facets of language generation in the limit. *arXiv preprint arXiv:2411.15364*.

Danescu-Niculescu-Mizil, C., West, R., Jurafsky, D., Leskovec, J., and Potts, C. 2013. No country for old members: User lifecycle and linguistic change in online communities. In *Proceedings of the 22nd international conference on World Wide Web*. 307–318.

Duetting, P., Mirrokni, V., Paes Leme, R., Xu, H., and Zuo, S. 2024. Mechanism design for large language models. In *Proceedings of the ACM Web Conference 2024*. 144–155.

Fish, S., Gonczarowski, Y. A., and Shorrer, R. I. 2024. Algorithmic collusion by large language models. *arXiv preprint arXiv:2404.00806 7.*

Gold, E. M. 1967. Language identification in the limit. *Information and control 10,* 5, 447–474.

Jagadeesan, M., Garg, N., and Steinhardt, J. 2023. Supply-side equilibria in recommender systems. *Advances in Neural Information Processing Systems 36*, 14597–14608.

Kalavasis, A., Mehrotra, A., and Velegkas, G. 2024a. Characterizations of language generation with breadth. *arXiv preprint arXiv:2412.18530*.

Kalavasis, A., Mehrotra, A., and Velegkas, G. 2024b. On the limits of language generation: Trade-offs between hallucination and mode collapse. *arXiv preprint arXiv:2411.09642*.

Kleinberg, J. and Mullainathan, S. 2024. Language generation in the limit. *Advances in Neural Information Processing Systems 37*, 66058–66079.

Kleinberg, J., Mullainathan, S., and Raghavan, M. 2024. The challenge of understanding what users want: Inconsistent preferences and engagement optimization. *Management science 70,* 9, 6336–6355.

Kleinberg, J. and Raghavan, M. 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences 118,* 22, e2018340118.

Kleinberg, J. and Wei, F. 2025. Density measures for language generation. *arXiv preprint arXiv:2504.14370*.

Li, J., Raman, V., and Tewari, A. 2024. Generation through the lens of learning theory. *arXiv preprint arXiv:2410.13714*.

Mladenov, M., Creager, E., Ben-Porat, O., Swersky, K., Zemel, R., and Boutilier, C. 2020. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *International Conference on Machine Learning*. PMLR, 6987–6998.

Raghavan, M. 2024. Competition and diversity in generative ai. *arXiv preprint arXiv:2412.08610*.

Raman, A. and Raman, V. 2025. Generation from noisy examples. *arXiv preprint arXiv:2501.04179*.

Taitler, B. and Ben-Porat, O. 2025. Selective response strategies for genai. *arXiv preprint arXiv:2502.00729*.

Yao, F., Li, C., Sankararaman, K. A., Liao, Y., Zhu, Y., Wang, Q., Wang, H., and Xu, H. 2023. Rethinking incentives in recommender systems: are monotone rewards always beneficial? *Advances in Neural Information Processing Systems 36*, 74582–74601.