Statement of Research

My research develops AI-driven, optimization-grounded methods for *computational marketing*, *online decision-making*, and *platform optimization*. I integrate tools from machine learning and operations research to study resource allocation, targeting, and influence in digital platforms, including influence maximization and fractional or discount incentives, online learning for subset selection and recommendation, and platform mechanisms and integrity. A related stream studies generative AI safety as strategic behavior and incentives, modeling attacker–defender interactions and incentive-compatible defenses; I also explore causal control in LLMs via latent or causal interventions to improve interpretability and controllability. Across these areas, I aim to combine theoretical rigor with deployable tools that inform data-driven decision-making in marketing analytics, platform design, and digital systems.

Algorithmic Marketing and Influence in Digital Platforms. Influence maximization addresses a fundamental decision problem in marketing and platform environments: given limited resources, which individuals in a network should we target to maximize adoption or engagement? The canonical problem is NP-hard (Domingos and Richardson, 2001), and while near-optimal greedy methods exist (Kempe et al., 2003; Leskovec et al., 2007), they rely on costly simulations at scale. To overcome these barriers, we proposed *Community-IM*, a community-aware divide-and-conquer framework that accelerates influence maximization by exploiting meso-scale structure, improving runtime while maintaining competitive influence spread on real networks (Umrawal and Aggarwal, 2023; Umrawal et al., 2023b).

We then generalized the model to *fractional influence maximization (FIM)*, also known as discount allocation, where incentives are divisible. We established structural links to classical IM and designed an efficient greedy algorithm with near-optimal guarantees (Umrawal et al., 2023a, 2026c). Building on this, we extended the framework to general marketing strategies with realistic, non-linear and heterogeneous adoption responses (Bhimaraju et al., 2024). Most recently, this line of work culminates in the *Budgeted Generalized Fractional Influence Maximization (BGFIM)* framework (Umrawal et al., 2026b), which incorporates heterogeneous user responsiveness and heterogeneous incentivization costs under budget constraints, together with a fast discrete-greedy algorithm that achieves a $(1 - 1/e)$ approximation guarantee while remaining scalable on large networks (manuscript in preparation for *Management Science*).

Building on Community-IM, we introduced *Community-IM++* (Robson and Umrawal, 2025), a scalable framework that explicitly models cross-community diffusion using a community-based diffusion-degree heuristic and progressive budgeting with lazy evaluation. On large real networks, Community-IM++ approaches near-greedy influence at substantially lower runtime while outperforming community-only baselines across budgets. Relatedly, we proposed a peer-pressure-based information diffusion model that dynamically scales a node's outgoing influence based on the cumulative pressure received from its activated neighbors, enabling richer modeling of network-dependent and reinforcement-based adoption dynamics (Stutsman et al., 2026), where adoption depends on cumulative social reinforcement; the extended version (Umrawal et al., 2026e) develops improved influence-maximization strategies and validates them using real-world diffusion traces (manuscript in preparation for *Management Science*).

Future work includes competitive multi-marketer targeting, fairness-aware allocation, and robustness to misinformation and strategic interference within evolving platform environments.

Online Learning for Targeting and Recommendation in Platforms. Many platform decisions reduce to selecting a subset under uncertainty, including seeding, ad placement, and recommendation. In the online setting, rewards are unknown and must be learned while optimizing long-run performance. We developed algorithms for combinatorial multi-armed bandits with non-linear feedback (Agarwal et al., 2021a, 2022, 2021b) and for submodular rewards under full-bandit feedback (Nie et al., 2022), introducing an Explore-Then-Commit Greedy approach with provable regret and strong empirical performance. We also developed adaptive top-n retained-set exploitation for weakly separated bandit settings (Anand and Umrawal, 2026) that, instead of committing to one leader, distributes exploitation across competitive actions, improving robustness in favorable regimes (manuscript in preparation for *Operations Research*).

Also, to reduce exploration cost in online influence maximization under full-bandit feedback, we designed the *Lazy Online Forward Algorithm (LOFA)* (Xu and Umrawal, 2026), which uses submodularity and lazy forward selection to prune candidates during exploration, achieving lower empirical regret on real-world networks. The ideas also extend to discount allocation and product recommendation.

Future work includes integrating fractional allocation with bandit learning, exploiting structure for sharper regret bounds, and incorporating knapsack-type constraints.

Platform Systems, Mechanisms, and Open-Source Tools. A core thrust of my work is translating models into deployable tools and platform mechanisms. We released open-source infrastructure for network diffusion modeling (Robson et al., 2024) with a maintained PyPI package, now pyOpenSci-verified for packaging standards, testing, and documentation. This infrastructure enables scalable simulation and analysis of diffusion processes for platform analytics. On the mechanism-design side, I study dynamic taxation and pricing for managing information overload and steering behavior, formalized as a sequential planner-sender game (Umrawal et al., 2026a,d), with implications for attention allocation and platform governance (manuscript in preparation for *Information Systems Research*).

Generative AI, Strategic Behavior, and Platform Integrity. I study generative AI systems as strategic environments shaped by incentives. We developed intent-hiding adversarial prompting and attacker–defender models, identifying structural vulnerabilities and guiding incentive-compatible defenses (Wu et al., 2026); the extended version strengthens the theory and empirical analysis (manuscript in preparation for *INFORMS Journal on Computing*). Relatedly, I also explore causal control in LLMs via latent interventions (Huang et al., 2025) to improve interpretability and controllability in AI systems.

Summary and Agenda. My research unifies computational marketing, online learning, and platform optimization with a focus on incentives, scalability, and impact:

- **Algorithmic marketing:** influence maximization; fractional and heterogeneous incentives; competitive and robust targeting; community-aware methods.
- **Online learning:** CMAB methods for recommendation and budgeted allocation with improved regret.
- **Platforms and integrity:** diffusion tools, attention mechanisms, and incentive-based generative AI safety and control.

The long-term vision is an integrated toolkit combining provable algorithms with deployable systems to support effective and data-driven decision-making in digital platforms and marketing environments.

References

Agarwal, M., Aggarwal, V., Quinn, C. J., and Umrawal, A. K. (2021a). Stochastic top-k subset bandits with linear space and non-linear feedback. In *Algorithmic Learning Theory*, pages 306–339. PMLR.

Agarwal, M., Aggarwal, V., Umrawal, A. K., and Quinn, C. (2021b). Dart: Adaptive accept reject algorithm for non-linear combinatorial bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6557–6565.

Agarwal, M., Aggarwal, V., Umrawal, A. K., and Quinn, C. J. (2022). Stochastic top k-subset bandits with linear space and non-linear feedback with applications to social influence maximization. *ACM/IMS Transactions on Data Science (TDS)*, 2(4):1–39.

Anand, A. and Umrawal, A. (2026). Adaptive top-retained-set exploitation in weakly separated bandit regimes. In *NeurIPS Workshop on Mathematical Foundations and Operational Integration of ML for Uncertainty-Aware Decision-Making*. Poster.

Bhimaraju, A., Robson, E. W., Varshney, L. R., and Umrawal, A. K. (2024). Fractional budget allocation for influence maximization under general marketing strategies. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 3627–3631.

Domingos, P. and Richardson, M. (2001). Mining the network value of customers. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 57–66. ACM.

Huang, Y., Chen, D., and Umrawal, A. K. (2025). Jam: Controllable and responsible text generation via causal reasoning and latent vector manipulation. *arXiv preprint arXiv:2502.20684*.

Kempe, D., Kleinberg, J., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146. ACM.

Leskovec, J., Krause, A., Guestrin, C., Faloutsos, C., VanBriesen, J., and Glance, N. (2007). Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 420–429.

Nie, G., Agarwal, M., Umrawal, A. K., Aggarwal, V., and Quinn, C. J. (2022). An explore-then-commit algorithm for submodular maximization under full-bandit feedback. In *Uncertainty in Artificial Intelligence*, pages 1541–1551. PMLR.

Robson, E., Reddy, D., and Umrawal, A. (2024). Cynetdiff: A python library for accelerated implementation of network diffusion models. *Proceedings of the VLDB Endowment*, 17(12):4409–4412.

Robson, E. W. and Umrawal, A. K. (2025). A community-aware framework for influence maximization with explicit accounting for inter-community influence. *arXiv preprint arXiv:2512.23973*.

Stutsman, C., Robson, E. W., and Umrawal, A. K. (2026). A pressure-based diffusion model for influence maximization on social networks. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 20, pages 2238–2251.

Umrawal, A., Robson, E., Zheng, A., Ly, E., and Varshney, L. (2026a). Regulating attention in digital platforms via dynamic taxation and strategic learning. In *HCOMP/CI Workshop on Cognitive Security for Human–AI Systems (CogSec)*. Talk.

Umrawal, A. K. and Aggarwal, V. (2023). Leveraging the community structure of a social network for maximizing the spread of influence. *ACM SIGMETRICS Performance Evaluation Review*, 50(4):17–19.

Umrawal, A. K., Aggarwal, V., and Quinn, C. J. (2023a). Fractional budget allocation for influence maximization. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 4327–4332. IEEE.

Umrawal, A. K., Bhimaraju, A., Robson, E. W., and Varshney, L. R. (2026b). Optimizing marketing budgets for social influence: A generalized fractional incentives framework. *arXiv preprint arXiv:2610.25823*.

Umrawal, A. K., Quinn, C. J., and Aggarwal, V. (2023b). A community-aware framework for social influence maximization. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 7(4):1253–1262.

Umrawal, A. K., Robson, E. W., Aggarwal, V., and Quinn, C. J. (2026c). Maximizing the spread of influence through a social network using partial incentives. *Journal of Artificial Intelligence Research (JAIR)*, 86(15):1–25.

Umrawal, A. K., Robson, E. W., Zheng, A., Ly, E., and Varshney, L. R. (2026d). Regulating attention in digital platforms via dynamic taxation and strategic learning. *Available at SSRN 6921899*.

Umrawal, A. K., Stutsman, C., Woods, C., and Robson, E. W. (2026e). How social reinforcement shapes word-of-mouth: A conviction-based model for marketing diffusion and network targeting. *arXiv preprint arXiv:2610.13890*.

Wu, X., Zhang, H., Umrawal, A., and Varshney, L. R. (2026). A game-theoretic analysis of attacks on large language models via compositional skills. In *Proceedings of the 43rd International Conference on Machine Learning*, pages 137513–137543. PMLR.

Xu, J. and Umrawal, A. K. (2026). Lofa: Online influence maximization under full-bandit feedback using lazy forward selection. *arXiv preprint arXiv:2601.00933*.