



Feature Based Dynamic Matching

YILUN CHEN, CUHK Shenzhen, China

YASH KANORIA, Columbia Business School and Amazon, USA

AKSHIT KUMAR, Columbia Business School, USA

WENXIN ZHANG, Columbia Business School, USA

Motivated by matching platforms that match agents in a centralized manner, we introduce a model of dynamic two-sided matching where both demand and supply are heterogeneous with many types and the pool of supply units is limited. We model heterogeneity on the two sides of the market by *i.i.d.* demand weight vectors and *i.i.d.* supply feature vectors, with possibly different distributions. The matching of a demand-supply pair generates a utility that depends on their weight and feature vectors. To reflect the realistic structure of a heterogeneous matching market while also avoid impossibility results, we consider various levels of assumptions (in particular, the spatial structure) on matching utilities and feature distributions. The goal of the centralized platform is to dynamically assign supply units to sequentially arriving demand units in order to maximize utility. Many popular heuristic policies are either sub-optimal (like the myopic policy) or computationally inefficient (like the certainty equivalent policy). We propose a forward-looking supply-aware policy dubbed Simulate-Optimize-Assign-Repeat (SOAR) that combines practicality and strong theoretical guarantee. Inspired by model predictive control (MPC), SOAR leverages the power of simulation to balance between producing immediate high match utility and preserving valuable supply for future demands. We use regret as our performance metric for matching policies, specifically the additive loss relative to the utility per match achievable in the continuum limit ($n \rightarrow \infty$). Under mild regularity assumptions on the offline matching instances, we prove that SOAR achieves the optimal regret scaling (up to a log factor). We further characterize the optimal regret scaling for interesting classes of problems with additional model structure, in particular, two classes of utility functions: (i) the “spatial utilities”, namely the negative p -th power of the Euclidean distance between the supply and demand vectors where $p \geq 1$; and (ii) the dot-product utility (equivalently $p = 2$ of (i)), and two classes of distributions: (i) both supply and demand distributions are smooth (a more stringent assumption) and (ii) supply and demand distributions are supported over compact sets (a mild assumption). En route to proving our guarantees we develop a novel framework for analyzing the performance of our SOAR policy which may be of wider applicability and independent interest. As a corollary of our techniques, we also resolve an open problem posed in Kanoria 2022.

A full version of the paper can be found at <https://ssrn.com/abstract=4451799>.

ACKNOWLEDGMENTS

Yash Kanoria and Akshit Kumar gratefully acknowledge the support of the National Science Foundation via grant CMMI-1653477.

ACM Reference Format:

Yilun Chen, Yash Kanoria, Akshit Kumar, and Wenxin Zhang. 2023. Feature Based Dynamic Matching. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23), July 9–12, 2023, London, United Kingdom*. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3580507.3597797>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

EC '23, July 9–12, 2023, London, United Kingdom

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0104-7/23/07.

<https://doi.org/10.1145/3580507.3597797>