

# Dueling algorithms

March 18 2011, by Larry Hardesty

---



Ankur Moitra, a PhD student in the Computer Science and Artificial Intelligence Laboratory. Photo: Patrick Gillooly

There's an old joke about two hikers on a trail, one wearing hiking boots and the other running shoes. "Why the running shoes?" the first hiker asks. "In case of bears," the second answers. The first hiker laughs and says, "Running shoes won't help you outrun a bear." "I don't need to beat the bear," the second hiker says. "I just need to beat you."

When computer scientists design algorithms, they're generally concerned with providing the best possible answer in the shortest amount of time. But for a web company in a competitive market, the best algorithm might just be the one that beats the other guy. At the Association for Computing Machinery's 43rd Annual Symposium on Theory of Computing in June, Ankur Moitra, a PhD student in the Computer Science and Artificial Intelligence Laboratory, and colleagues will

present a new mathematical framework for analyzing such “dueling algorithms.” The work could help answer questions such as when competition between web services helps meet social ends and when it undermines them, and it could also prove useful in the social sciences.

To understand the idea of dueling algorithms, Moitra says, consider a [search engine](#) ranking the results of a query. The words of the query could be subject to multiple interpretations: “Cambridge,” for instance, could refer to several different cities. Suppose that in using a given search term, 40 percent of people mean one thing, 30 percent a second and 30 percent a third. A computer scientist designing a search algorithm in the abstract would probably prioritize the most likely interpretation. But in a commercial market, if that’s what the leading search engine does, a competitor might do well to prioritize the other two. Forty percent of its customers will be disappointed, but the other 60 percent will prefer the results they get to those provided by the market leader.

Of course, the leader is unlikely to sit idle as the upstart steals its business. So it might evolve some hybrid ranking strategy, sprinkling some of the less likely interpretations among its top results. That, in turn, would prompt the competitors to revise their strategies, which would prompt another revision from the leader, and so on. Ultimately, the theory goes, the competitors in the market will converge on what economists call an equilibrium, a state in which no competitor has any incentive to change its strategy unilaterally. Any given contest could have many equilibria; which one emerges depends on how the contestants’ strategies evolve over time.

## Balance of power

Moitra and his colleagues at Northwestern University, the University of Toronto, the University of Pennsylvania, Israel’s Technion and Microsoft Research have developed methods for finding equilibrium strategies

much more efficiently than was previously possible, at least for two-player, winner-take-all contests. Even these simple contests, however, yield enormously complex calculations. Finding equilibria generally requires comparing all possible strategies against each other. In the search engine example, that would mean comparing every ordering of search results against every other ordering. As the number of results gets larger, the complexity of the comparison increases exponentially.

For several different types of contest, however, Moitra and his colleagues found ways to represent the results of different strategies probabilistically. That makes it much easier to calculate equilibria, but the strategies found to be in equilibrium are defined only by their statistics. So the researchers also had to provide computational methods for finding specific sets of strategies that match the statistical profiles of the equilibrium calculation.

In their paper, the researchers consider several cases of dueling algorithms. Among them are two computers searching for an item in a database, two companies trying to hire employees from the same pool of applicants, and two racers trying to plot routes across a town with erratic traffic patterns. In each case they find that an equilibrium strategy for beating an opponent may not be a good strategy in the abstract: the company may not end up with the best possible employees, for instance, and the computer may not find the item in the database as efficiently as it could.

According to Eva Tardos, Jacob Gould Schurman Professor of Computer Science at Cornell University, the researchers' paper "is more the beginning of research than a definitive result or end product." The researchers' models make several simplifying assumptions — including the number of competitors — that make the math easier but limit their applicability. Nonetheless, "just raising the questions is an important step forward," Tardos says. "With traditional optimization, you just want the

[algorithm](#) to be as good as possible, versus wanting to beat the opponent. Clearly, there's a tension between these goals. I think they are setting a research agenda to understand this tension.”

---

*This story is republished courtesy of MIT News ([web.mit.edu/newsoffice/](http://web.mit.edu/newsoffice/)), a popular site that covers news about MIT research, innovation and teaching.*

**More information:** [arxiv.org/pdf/1101.2883v1](https://arxiv.org/pdf/1101.2883v1)

Provided by Massachusetts Institute of Technology

Citation: Dueling algorithms (2011, March 18) retrieved 20 March 2024 from <https://phys.org/news/2011-03-dueling-algorithms.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--