

Expert discusses the 'horrendously complex science' behind online advertising research

May 21 2021, by Steve Hawk



Credit: CC0 Public Domain

As an expert in the algorithmic engines that drive online advertising, Harikesh Nair loves to explore that rarified zone where deep scholastic research and real-world business needs overlap.

Which is one reason why, in 2017, he took a leave of absence from his tenured professorship at Stanford Graduate School of Business and joined the engineering team at JD.com, one of China's largest e-commerce platforms. He spent two years shuttling between Silicon

Valley and Beijing, overseeing a team of engineers working in the trenches to convert dense mathematical concepts into tangible online solutions.

When he returned to Stanford GSB in 2019, Nair helped found the school's Computational Marketing Lab, which brings together faculty, students, and practitioners to identify [best practices](#)—both in terms of profitability and societal benefit—for online marketers.

"It used to be that only scientists and academics and engineers were thinking about [data science](#), but now it has become a C-level issue," he says. "The CTO, the CMO, the CEO—everybody's talking about it. And if they're not, the board will ask about it."

Insights by Stanford Business recently sat down with Nair to hear his thoughts about the current state of [online advertising](#), the challenges of running experiments when real-time commerce is measured in millionths of a second, and what kinds of digital ads annoy him.

Why exactly did you take a leave of absence to work at JD.com?

Since I received my Ph.D. in 2005, the practice of marketing in industry has become very computational. It's quickly on its way to becoming almost entirely a data-driven field. If you think about the companies where amazing marketing was originally invented—places like Clorox and other packaged brands—that's no longer the case. These days, all of the significant innovations in marketing and advertising and pricing are coming out of tech companies like Google and Facebook and Amazon. That was a trend I wanted to be part of.

Why JD.com?

I'm interested in retail and e-commerce, and in that world the gravity has shifted to Asia. Silicon Valley is still an innovation hub, but the big growth in online retail is happening in Asia. The middle class there is expanding, so people are making more choices based on the quality, or perceived quality, of goods. It's aspirational consumption. At some point I got connected to some senior folks from JD, and we talked about research and collaboration. There was an opportunity to spend some time with them and dive much deeper. So I thought, "Hey, this is a good opportunity: tech and Asia."

Aside from the different sizes of the markets—billions of potential customers in Asia vs. hundreds of millions in the U.S.—what makes online marketing in Asia different from here?

In the U.S., we're still a desktop nation. People still use Chrome or Safari on their laptops or home computers to shop at Amazon.com. In Asia, and especially in China, very few people use that stuff anymore. Everything's done through the [mobile phone](#); 80% of shopping traffic comes through mobile apps, and a lot of consumers are in the generation that skipped laptops entirely and went directly to mobile phones. That changes how commerce works.

How so?

It's all about discoverability. With [mobile apps](#), there's limited real estate: Nobody browses beyond the first three pages. If you have 8 listings per page, that's 24 products. So you have to surface 24 items out of 300 million possible product options, and you have to personalize the rankings by finding a good match between the user's needs and what's available. All of that has to be done really fast, within 200 microseconds [0.0002 seconds], and it has to be executed flawlessly. It's a horrendously complex science problem.

How does a consumer's demographic characteristics weigh into that problem?

Actually, age and gender and things like that are not that informative. What they bought recently and what they searched for is what matters.

Beyond the overwhelming use of mobile in China, are there other major differences between Chinese and American consumers in the way they interact with e-commerce sites?

In China, the social aspect of consumption is a major factor. Their society is more communal than ours, so consumption is more embedded in their social networks. Rather than go to a site and search for something and then click and buy and leave, which is typical American consumption behavior, people in China share products on each other's social feeds, and by doing so they can often get a price reduction. It is also a value-conscious society.

Another important difference is that the United States has three big companies—Facebook, Amazon, and Google—that don't interact. I think of them as walled gardens. Facebook is a social platform with lots of advertising but very little in the way of transactions. Amazon is where the commerce happens, but it doesn't have much of an ad business—although that's changing. And Google is where search occurs, but it doesn't have much of a transaction or social business. But in China, you have Tencent, which does all three: search and social and commerce. It's well merged. Same with Alibaba. They both have their own ecosystems. And within each ecosystem, it's all seamlessly integrated, so you flow smoothly from the social interaction to the search to the transaction to the payment. Consumers respond well to that.

Given that marketing is almost entirely data-driven

these days, is there any art to it anymore?

I don't think the art has gone away, but it has shifted. A good manager needs to understand how to leverage algorithms and also understand what the data can do and cannot do. They need to know when to base a decision on the data and when to base it on their experience: What human capital is required to convert the data into something actionable?

We recently ran a video about your work in which you said, "Just because I bought something after I see an ad for it doesn't mean the ad caused me to buy it." Can you talk more about that?

That issue is a huge one and it's very, very complicated in advertising markets. The number one problem facing any brand that's spending money on ads is measurability: They want to know not just how many people saw their ad, but also whether they're getting a return in terms of sales. Proving that is very hard, because I can't just count up the people who saw an ad and then bought the product, because maybe they were already planning to buy it. Fifteen years ago, a platform probably could have used that kind of flawed analytics, but brands aren't buying that anymore. They want proof that their ads are causing incremental increases in sales. Incrementality is the number one issue in the ad business right now—proving that properly.

How do you prove it?

Typically by running experiments at scale. The experimentation is very complex, because it's impossible to design and run a separate experiment every time an individual advertiser wants that measurement. When I was at JD, for instance, we had hundreds of thousands of advertisers. Most ads are sold by auction, so you have to encode the experiments into the auctions. In order to show somebody an ad, you've got to win the

auction, and if you want to make sure certain people don't see the ad—for the sake of the experiment—you have to have cases where the advertiser purposefully loses the auction. Both circumstances have economic consequences, because losing the auction may cost you the sale and winning the auction requires you to pay a bid.

Also, the experiment cannot slow down the system; you can't have latency. So you've got to build out the infrastructure, get the statistics encoded, and then make it an external-facing, self-serve mechanism that the advertiser can turn on and off. All of this has to be running in the background—the complexity has to be hidden. And then the results have to be communicated through a simple interface visualized into meaningful insights for people who don't really know that much about statistics.

What lessons that you've learned have blown up some of the assumptions that marketers used to have?

I think most researchers who work in this space would agree that, historically, the effect of advertising has been oversold—that the true causal effects of ads are actually much smaller than what some in the industry have claimed. Another thing we've learned is that certain kinds of ads can be very effective. For example, search advertising—the ads that appear in response to your online searches—is actually quite powerful. Data has shown that consumers actually like that. When we take that away, in properly measured experiments, we find that users stop coming to the platform. They actually prefer platforms that have those informative ads. My colleague Navdeep Sahni here at Stanford GSB has run some really beautiful experiments showing that.

One last question: How has your research changed the way that you personally respond to ads?

I'm annoyed by annoying ads now, because I know there's bad science behind them. I'm OK with data being collected and things like that, but my tolerance for lack of transparency is very low. And of course, whenever an ad appears in my feed now, I always wonder, "What's the auction system that caused this result? What's the algorithm behind it?" That's the fun stuff.

Provided by Stanford University

Citation: Expert discusses the 'horrendously complex science' behind online advertising research (2021, May 21) retrieved 20 June 2024 from <https://techxplore.com/news/2021-05-expert-discusses-horrendously-complex-science.html>

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.