FORTUNE

How marketers are increasingly using A.I. to persuade you to buy

BY
DAVID Z. MORRIS
January 31, 2020 6:30 AM EST

Assaf Baciu, co-founder of Persado, a startup that uses artificial intelligence to help companies write ads, turns his back to me.

"If I'm not paying attention to your reactions," he says. "It's as if I'm not listening to you. That wouldn't be very human, would it?"

Baciu's point is that truly human marketing means constantly listening to what individual consumers really want, and appealing to their desires. But using actual humans to do that would be far too expensive.

His company creates software that helps advertisers listen to their millions of customers, and evaluate whether those customers pay attention to and are persuaded by the marketing they see. The technology identifies what tactics work best for particular kinds of consumers, then suggests messaging tailored for them.

Persado's work is helping to give rise to data-driven advertising, and the increasing use of artificial intelligence to help make it more effective. Marketers hope that technology like Persado's can push the trend further by taking the guesswork out of ad writing, a task that still relies largely on the instincts of mere humans.

Although data-driven targeting has revolutionized advertising, there is growing concern about its social and psychological impacts. It also faces growing criticism of its privacy implications, earning it the ignominious distinction of being lumped in with what is increasingly called "surveillance capitalism."

Since its founding in 2012, Persado's clients have included the likes of Dell, Staples, and American Express. In July, Persado announced a five-year deal with J.P. Morgan Chase to improve both its advertising and its automated customer service 'chatbots.'

Persado has a lot of competition in the automated ad copywriting field. Other startups with a similar focus include Motiva AI and Phrasee.

A.I.-driven copywriting leverages the same combination of user data and machine learning that Facebook and Google use to deliver individually-targeted ads. But unlike the tech giants and their fondness for constant surveillance, Persado says its tool doesn't actively track or profile individual users. Instead, it uses broad demographic data and ad-performance results from clients to guide its ad-polishing machines, while delivering ads to viewers is left to each client.

But A.I. copywriting could be seen as amping up targeted advertising, and its pitfalls, to a new level. Instead of just delivering ads to the right people, Persado makes it possible for its clients to tailor ads themselves in a highly individualized way.

And because the software is also highly automated, it can be a good listener – and therefore very persuasive – at a massive scale.

The influence machine

Persado is not quite ready to replace ad copywriters wholesale – the software creates and tests variations of existing ads, rather than creating them from scratch. "There's no machine that writes on its own," says Baciu. "It all comes from humans, originally."

Instead, Persado works by breaking human-crafted ads into key elements such as layout, imagery, and language. Using both customer response data and a database of more than a million common marketing phrases that have been categorized and labeled by Persado's human staffers, the software then automatically generates 16 alternate versions of the message. They can vary in everything from specific phrasing, to color scheme, to fundamental emotional appeal.
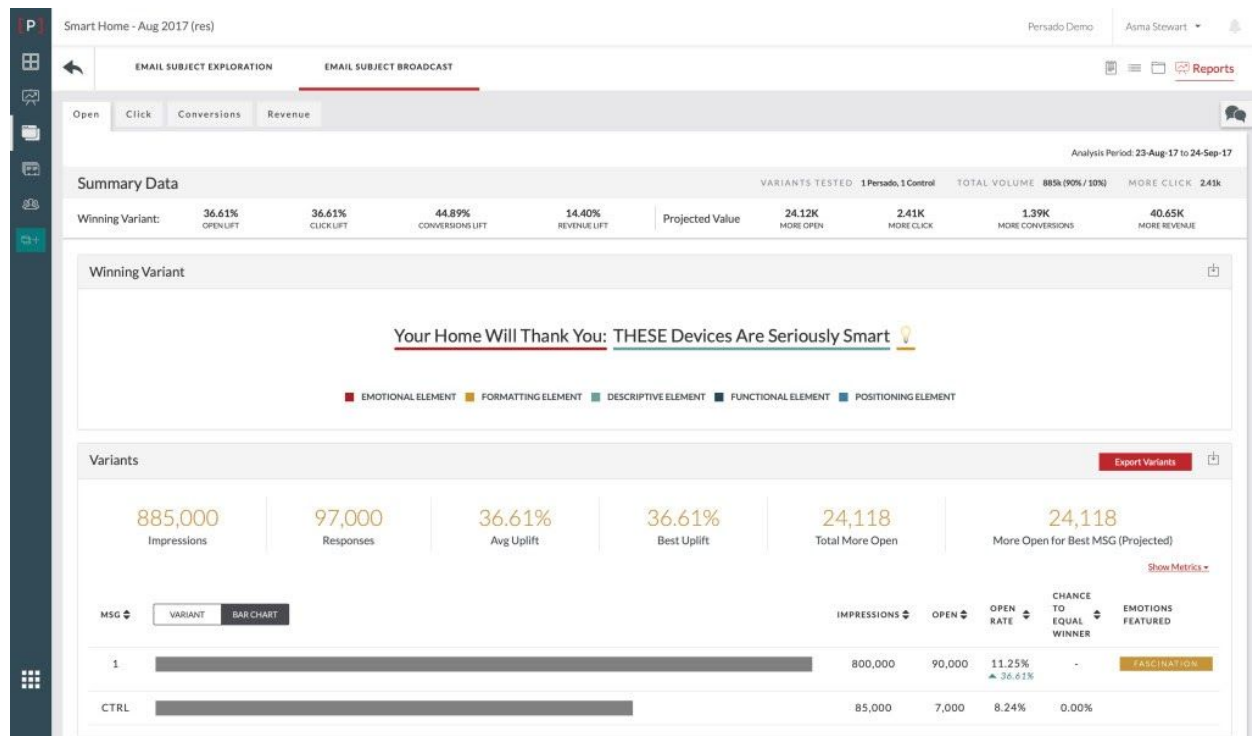
These variations are then shown to test audiences and measured for effectiveness, in a process broadly similar to the A/B testing often performed by marketers – but far more detailed. According to Persado, the tests of those 16 ad variations are structured in a way that provides insight into the effectiveness of thousands of possible variations and combinations.

"It's as if you're asking customers which ideas resonate with them and which don't," says Baciu. "It's like a massive survey that's continuously learning."

The most effective variants are shown to Persado's clients for review before they choose which versions to show to a larger audience. Humans must approve the language in every message before sending them out.

The results can seem like the outcome of nuanced psychology. For instance, after starting with ad for an electronics sale that played on a sense of urgency ("Over before you know it"), Persado found it was more effective to focus on exclusivity ("You're on our list"). But like any A.I., Persado's algorithms lack any real

understanding of the words being shuffled around. "The machine is creative in the sense that it can be exhaustive," says Baciu.

Persado's user dashboard, showing how the artificial intelligence can categorize elements of a marketing message, then suggest variations and test their effectiveness.
PERSADO

Persado's work may be effective partly because, in a world where people see an average of 5,000 commercials, sponsored tweets, and subway placards daily, advertising has become increasingly self-defeating. "We have our filtering hats on 24/7 and we're just filtering out messages," says Ari Lightman, professor of digital media and marketing at Carnegie Mellon University. Knowing exactly which psychological buttons to push can help defeat those heightened defenses.

"Neural engineering"

As appealing as bypassing that noise is for businesses, the idea of persuasive technology is also more than a little foreboding. "I think there are definitely ethical concerns," says William Wang, director of the Center for Responsible Machine Learning at the University of California at Santa Barbara. "You have to be very careful when you are trying to persuade someone using A.I."

That's in part because these systems have simple goals, and no inherent ethical controls. Software like Persado's is designed to trigger the behavior an advertiser wants, and all other considerations are secondary.

Futurist Gray Scott describes work like Persado's as "neural engineering," and sketches some bleak possibilities. "Behavioral A.I. doesn't care if your grandmother can't pay her bills," says Scott. "She's hitting the buy button, even if she knows she shouldn't."

"Is it okay to use behavioral algorithms on children?" Scott asks. "On people making less than $17,000 a year? There are financial inequalities that are going to come from some of these algorithms."

A growing number of experts are focusing on the ethics of A.I., and a set of common standards is emerging. Communication A.I. should identify itself as a robot; companies should have ethics boards that review the social impact of what they're building; and, perhaps most importantly, A.I. should be able to explain why it makes certain decisions or suggestions, making it easier for humans to evaluate potential impacts.

Back in Persado's office, Baciu himself lays out an example of targeting that hints at exactly these risks. A financial services company's elite customers, he posits, may be most influenced by ads equating a product with 'proactivity,' greater control, and individual insight. New prospects, by contrast, are more likely to respond to a pitch

based on 'security,' or the trust they put in an institution. In other words, Persado can help companies figure out how to best influence less-savvy, less-confident customers.

Persado doesn't handle political ads, a decision it reached well before the likes of Google and Twitter clamped down on them. And Baciu says Persado only works with advertisers it deems "ethical" – though that's a rather slippery designation.

Persado also requires manual human approval of messages generated by its system. And as part of that review, Persado's dashboard breaks down why a particular message is more effective – whether because of stronger language, a different theme, or a better image.

That's particularly important because Baciu says about 60% of the improvement Persado makes to messages' effectiveness comes from honing appeals to viewers' emotions. And there is increasing evidence that automated systems designed to elicit human responses, such as increasing user engagement, tend to foster negative emotions.

It's easy to imagine an A.I.-driven ad robot discovering that self-loathing, fear, or resentment can be effective sales tactics. Deployed at mass scale, such messages could cause widespread harm.

Baciu claims that human oversight paired with Persado's dashboard actually helps steer many clients away from negative selling tactics. "In most cases, we actually move them from the very tactical messages," leveraging feelings like urgency, anxiety, or guilt, towards "something much more exciting – achievement, or gratitude, or intimacy, [emotions] that they don't even know they could use."

Even people like Wang, the U.C. Santa Barbara director, acknowledge that there may be social benefits to A.I.-driven communication. For instance, it can help provide consistent quality in customer service, either through fully automated chatbots, or by providing consistent scripts for human representatives.

But Baciu acknowledges that software similar to Persado's will eventually be more widely available, so leaning on the ethics of individual advertisers won't work for long. And there's never been a shortage of sales teams willing to use misleading or manipulative messaging, even on behalf of 'respectable' brands like Wells Fargo.

There are some legal controls that apply to all advertising, such as those barring outright false advertising. But there remain very few U.S. federal regulations of A.I., and none specific to their use in influencing human behavior.

So even Baciu strikes a cautious note, saying "A.I. should be somewhat supervised, if not regulated."