

A young man and woman are standing in front of a wall made of horizontal wooden slats. The man, on the left, is wearing a light blue t-shirt and a backpack, and is holding a smartphone. The woman, on the right, is wearing a grey crop top, blue jeans, and red-tinted sunglasses, and is also wearing a backpack. They are both smiling and looking at the phone. The text 'UNDERSTANDING LIQUIDITY' is overlaid on the left side of the image.

UNDERSTANDING LIQUIDITY

How machine learning helps media teams work smarter

facebook IQ

WHAT IS LIQUIDITY AND HOW DOES IT WORK?

Back in 2006, NASA engineers faced a dilemma as aircraft became increasingly automated, leaving pilots to input commands and manage automation sequences. The team posed the following question: How do we find a balance between using increasingly powerful technologies and retaining authority, with clear roles between humans and automation?

The answer was the H-metaphor,¹ a model for interaction with intelligent machines akin to a horse and its rider. According to the H-metaphor, humans chart the course, then rely on machines to manage real-time calibrations—much like a rider who trusts his horse to negotiate the terrain.

In a sense, this is where we are with machine learning and media planning. Media planners accustomed to crafting masterful plans have long been limited to setting very narrow parameters for platform buys. Meanwhile, buyers accustomed to buying and serving ads via automated systems still spend much of their time inputting, tweaking and honing campaigns. With myriad ever-changing data points to consider—for example, audience targeting, market fluctuations, creative executions and placement options—the calculus can become time-consuming.





In such a fluid marketplace, machine learning helps to identify the most valuable impressions. **When every dollar is allowed to flow to the most valuable impression, we call this condition “liquidity,” and it is made possible when humans take their hands off the controls and allow the system to read the terrain.**

Machine learning uses algorithms to determine trends and relationships within a dataset, which can be used for explaining observed effects or predicting future outcomes in real time. Systems based on machine learning are able to evaluate opportunities faster than any human being possibly could. In advertising, machine learning allows media teams to automate the more tedious parts of their jobs, like deciding which calls to action work best for which audience, where to serve the ads and when.

It’s understandable that marketers haven’t yet taken full advantage of this new technology. Algorithms can feel opaque, and some people worry that machine learning—and artificial intelligence in general—could make their roles disappear. But according to McKinsey,² the vast majority of positions can’t be fully automated using current technologies. By contrast, about a third of most tasks that employees perform could be automated, improving productivity and freeing employees to focus on higher impact work.


In this paper, we’ll explore just how machine learning enables liquidity within automated campaign systems that allow buying across multiple locations and systems. We’ll also explore how both planning and buying teams can use this technology to run more cost-effective campaigns.

UNDERSTANDING HOW MACHINE LEARNING WORKS IN AUTOMATED SYSTEMS

Embedding machine learning within automated systems is a natural evolution of today's automated bidding systems, which the digital media industry has been using for more than a decade. But while traditional exchanges were optimized solely for price and volume, machine learning allows for more complicated campaign optimizations.

Traditional ad exchanges require buyers to bid on specific inventory shown to a predetermined audience and against a cost metric. So, if a planner wanted to show a 250 x 250 banner ad to mothers of teenagers and pay no more than a \$10 cost per thousand (CPM), a traditional exchange could do that job. But when it's time to evaluate performance, it might be unclear what about that ad performed—audience, placement or creative? To find out, a media team might have to execute, run and analyze a dozen experiments before deciding on the best buy.

Machine learning makes this entire process much easier, as it supports systems that automate the process of finding the best buy. Machine learning models ingest behavioral data that consumers consent to provide to advertising platforms, website cookies or advertisers themselves. This data, which is anonymized and aggregated, is used to predict which groups of people will respond to an ad. Then, these models can select which audiences (or groups of people) to show the ad to, and when. The more information provided to the models, the better the predictions, and the better the system is able to identify and capture the impressions that are most valuable for meeting the advertiser's goal.



LEAVING MACHINES ROOM TO RUN

Media planners accustomed to creating very detailed maps for client campaigns might be tempted to specify who to target, where and with what creative. But planners can also use machine learning to surface insights and spark new ideas. To do so, they need to allow the system to explore as many opportunities to find value as possible, by setting broad parameters whenever they can. After the campaign runs, advertisers can then use the system's reporting tools to retrieve insights. This allows learnings from campaign performance to be surfaced without having to run many experiments manually.³

In many cases, machine learning within automated systems can bring planners closer to success. Campaigns aiming for wide reach can run across large audiences and platforms, allowing more choice and accessing valuable and efficient impressions at scale. When aiming for specific outcomes, marketers who choose broad audience parameters and an agnostic approach to platforms can produce more efficient conversions.

In every case, allowing machine learning algorithms to process as much data as possible improves liquidity. By removing restrictions on a campaign, media teams can test more opportunities simultaneously—including

creative iterations, placements, audiences and more—and ensure budgets are applied to find the right people and produce the most cost-efficient results.

The concept of cost per incremental outcome is useful in explaining exactly why this happens. Cost per incremental outcome estimates how much it costs an advertiser to change the attitude or behavior of a given person with an advertising campaign, and captures the relationship between campaign spend, ad effectiveness and reach and frequency.

**Cost per incremental outcome =
Spend/(Reach x Lift)**

If done right, we expect increased liquidity to result in increased reach, lift or a combination of the two, as the system finds more or better opportunities to show the ad for the same cost. Combined with constant campaign spend, this often drives lower cost per outcome.

In general, media teams can think about increasing liquidity on four main dimensions of a campaign: placement, audience, budget and creative. While different advertising platforms present different choices during campaign setup, these four dimensions are often nodes where planners and buyers must

choose whether to restrict opportunities considered for delivery during the campaign. Setting as few restrictions as possible, thereby increasing liquidity on as many dimensions as possible, should confer the greatest benefit.

HOW SHOULD I USE REPORTING?

While reporting tools often show breakdowns by creative, demographics, placement or other dimensions, planners should evaluate the success of the campaign by looking at the results in aggregate. It may be tempting to look at results and decide to exclude the worst-performing creative or placement in the next campaign. However, this is likely to result in worse campaign performance. An automated system can see the distribution of prices in real time, while a planner cannot, so it is able, for example, to shift delivery away from one placement if it is starting to experience rapid price increases. As a result, even if one segment appears to perform worse, delivering to this segment was likely more cost-efficient than the next-best inventory from other segments. This issue is called the breakdown effect—[learn more](#) with Facebook for Business.

4 MAIN DIMENSIONS OF LIQUIDITY

PLACEMENT

If the advertising platform includes multiple surfaces, placement liquidity can be increased by opting in to as many surfaces as possible.

AUDIENCE

Audience liquidity can be increased by targeting the broadest possible audience.

BUDGET

Budget liquidity can be increased by removing restrictions on where the campaign budget can be spent. For example, setting the budget at the campaign level instead of the ad level is one way to increase budget liquidity.

CREATIVE

Creative liquidity can be increased by allowing the automated system to test and choose the best-performing creative.



Teams that use machine learning-enabled systems can reap the benefits of increased liquidity, namely more efficient setup and more efficient campaigns.

This isn't to suggest that an automated system is a magic bullet. For it to succeed, media buyers must first identify a business goal to use in order to assess which opportunities are most valuable, then pass this information to the system by setting an optimization objective. The system will then listen for the appropriate signal that the goal is being achieved, and move budget and impressions to different ad surfaces or audiences that give the maximum effect.

What's more, a system is only as good as the network it runs across. Inventory quality, fraud prevention, brand safety protections and more should all be in place before you opt into any network's automated system. This is already part of any marketer's due diligence, but concerns become magnified within automated systems. If, for example, a system is optimized toward clicks, accidental clicks could lead to high bounce rates from the advertiser's website, skewing results.



CHOOSING SIGNALS


Marketers who entrust their campaigns to machine learning aren't passengers on the road to success—they're much-needed captains who understand that reaching a business objective is more complicated than ever. Their insights and contributions are invaluable in steering automated systems toward the right outcome.

A critical step in campaign success is choosing the optimization goal and associated signal that best aligns with your business goal. For example, an online retailer who cares about purchases should set a goal (objective) of conversions and optimize for a purchase event. When customers engage with an ad in a specific way—like adding an item to their cart—they're also sending a signal of purchase intent.

For marketers with direct-response goals, these signals often align closely with bottom-line objectives. App downloads, lead generation and online sales are all clear direct-response signals. By choosing an optimization signal that relates to your business goals, you are opting for a machine-learning model aligned with what you want the campaign to achieve. Conversely, selecting a signal not directly related to business goals may produce unexpected results. For example, an advertiser who chooses to optimize for link clicks might see a lower cost per click, but post-click conversions may suffer.

WHAT'S A SIGNAL?

A signal is a piece of consumer behavioral data that a machine-learning model can use to make predictions about future behavior. For example, if consumers add specific items to their carts, this may indicate they are likely to purchase in the future. Signals can come from a variety of sources such as through pixels (which can be set up for a range of events), software development kit (SDK) data that indicates app installs or data for offline events like in-store sales.



By contrast, marketers running brand campaigns may not have access to a signal closely aligned with bottom-line objectives. Commonly used signals like video views might not track directly to true business goals such as raising sales. Further, there may be a trade-off between different metrics. For example, if a campaign is optimized on video views, an automated campaign that is more liquid will deliver the cheapest cost per view—but may not deliver the highest reach.

These secondary effects are not specific to automated campaigns but are a potential result of using higher-funnel metrics as a stand-in when lower-funnel metrics are not available.

To determine which signal works best, brand marketers should use lift testing to evaluate how well the proxy metric can serve their business goals. Advertisers can also run a split test to see how well different signals work across a portfolio of metrics when using automated systems.

Media buyers who understand what information an automated system needs add value by finding the signals that best deliver against their client's goals.



HOW CAN I TEST MY CAMPAIGN?

Like an A/B test, split testing allows you to create multiple ad sets and test them against each other to see which strategies produce the best results. Split testing is good for comparing performance to decide between two strategies. Lift testing, on the other hand, shows the impact of an ad or campaign on business goals, compared to a world in which the campaign was not run. [Learn more](#) about the difference—and how to set up an effective split test—with online courses through Facebook Blueprint.


WHAT IT MEANS FOR MARKETERS

Media planners and buyers play an essential role in crafting and executing platform campaigns. But even the most masterful can't out-calculate a machine. And why would they want to? Machine learning-driven systems save time and effort when setting up campaigns, help maximize liquidity for the most efficient buy and uncover fresh insights that media teams can deliver to clients.

REMEMBER THAT BROADER IS (USUALLY) BETTER

Planners are trusted by the clients for their vision and understanding of consumer audiences. They also add value by turning up unexpected insights. Automated systems underpinned by machine learning can do just that by finding undiscovered pockets of efficiency that help clients reach goals faster and more cheaply.

If you are targeting a necessarily narrow audience or have other marketing guardrails, it's important to understand that automated systems may not have enough data to achieve the same results they otherwise might. In these cases, the ability to navigate the media landscape yourself is likely worth the additional cost.



WHAT IT MEANS FOR MARKETERS

OPT IN TO AUTOMATED SYSTEMS

How complicated are your objectives? If your goals are unorthodox, you care about multiple brand metrics or you are courting a niche audience, you may not be able to provide the system with enough information to realize your goals.

Further, how much control do you need? A manual setup should give you control over something that's not available in the automated system, like capping shares of delivery to a certain placement. If you don't need to set those kinds of granular parameters, manual optimization could lead to a less-efficient campaign—choose automated instead.

GUIDE THE SYSTEM

Is there a good signal that you, the advertiser, can provide to the system? This might be a true business goal, like lead generation or app downloads. If you have a strong signal, opt into automation without additional testing.

If not, how open are you to a trade-off between different campaign metrics? When you stray from your bottom-line goals, you increase the chances that your outcome may not drive your goals. If you have limited tolerance, try lift testing to see which strategy offers the best value, or split testing to get a quick read on how different optimizations affect a portfolio of metrics.



FOOTNOTES

¹ Kenneth H. Goodrich et al., “Application of the H-Mode, a Design and Interaction Concept for Highly Automated Vehicles, to Aircraft,” 25th Digital Avionics Systems Conference, October 15, 2006.
<https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20060051774.pdf>

² James Manyika et al., “Jobs Lost, Jobs Gained: What the Future of Work Will Mean for Jobs, Skills and Wages,” McKinsey Global Institute, 2017.
<https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages>

³ The level of information available in reporting tools varies by platform. In a Facebook-specific context, advertisers can use Ads Manager and Facebook Analytics for granular information about campaign performance.

facebook IQ