DIGITAL ATTRIBUTION PRIMER 2.0

August 2016

© 2016 Interactive Advertising Bureau Attribution Primer 2.0

This document has been developed by the IAB Advanced Attribution Working Group.

The IAB Advanced Attribution Working Group is part of the IAB Performance Committee, which is dedicated to providing industry perspectives on accountable media through best practices on campaign development, digital insights, measurement, and attribution—beyond the last click. The committee examines matters across multiple platforms and devices.

The Advanced Attribution working group aims to address the digital advertising community's challenges around adoption of campaign measurement approaches beyond first/last touch methodologies, specifically: 1) the lack of education around the limitations of first/last touch, 2) awareness of alternatives to first/last touch models, 3) complexity of implementation of fractional attribution tools and counting methodologies, and 4) lack of clarity around how best to applying advanced attribution data within the reporting and optimization process.

IAB Advanced Attribution Working Group Roster

Abakus
Adform
AdRoll
Affinity X
Amobee
AT&T AdWorks
Conversant Media
Convertro
Criteo
DashBid
DataXu
Demand Media
Dun & Bradstreet
Electronic Arts
Ernst & Young
Factual
Flipboard
Flite
Foursquare
GasBuddy.com
Google
GumGum
HookLogic
InMobi
Integral Ad Science
Kochava Inc.
M&C Saatchi Mobile
Madison Logic
Marchex
Matomy Media
Medialets
MediaMath
NinthDecimal
Oracle
Phluid Media
Pinterest
PubMatic
PulsePoint
Qualia Media
Quantcast
Rakuten Marketing
Research Now
ResponsiveAds
Rocket Fuel
Sizmek
StartApp
Tavant Technologies
The Trade Desk
The Weather Company, an IBM Business
Turn Inc.
Twitter
WebbMason, Inc.
WebMD
xAD
YuMe
Zemanta

About IAB
The Interactive Advertising Bureau (IAB) empowers the media and marketing industries to thrive in the digital economy. It is comprised of more than 650 leading media and technology companies that are responsible for selling, delivering, and optimizing digital advertising or marketing campaigns. Together, they account for 86 percent of online advertising in the United States. Working with its member companies, IAB develops technical standards and best practices and fields critical research on interactive advertising, while also educating brands, agencies, and the wider business community on the importance of digital marketing. The organization is committed to professional development and elevating the knowledge, skills, expertise, and diversity of the workforce across the industry. Through the work of its public policy office in Washington, D.C., IAB advocates for its members and promotes the value of the interactive advertising industry to legislators and policymakers. Founded in 1996, IAB is headquartered in New York City and has a West Coast office in San Francisco.

IAB Contact
Benjamin Dick
Director, Industry Initiatives
Programmatic / Data / Performance
ben@iab.com

© 2016 Interactive Advertising Bureau Attribution Primer 2.0

Table of Contents

1 Attribution Definitions

1.1 Attribution

Attribution is the process of identifying a set of user actions ("events") across screens and touch points that contribute in some manner to a desired outcome, and then assigning value to each of these events.

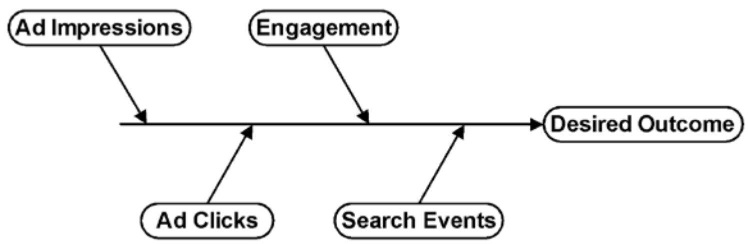

In digital advertising, attribution measurement is most widely performed by media buyers trying to understand which aspects of digital media campaigns are contributing most to campaign performance. Digital attribution is done at a user-specific level, through efforts to assign a consistent user identifier across all analyzed events. This is opposed to traditional media performance analysis, where ROI is generally understood at the macro user-group level because there is no consistent user identifier available (see Marketing Mix Modeling, section 4.1).

The focus of this document is user-specific attribution, but it will also reference user-group attribution (i.e., Marketing Mix Modelling) to highlight the challenges of integrating data across media.

1.2 Attribution Ecosystems

Attribution is a necessary component of digital advertising. Without a basic understanding of which ads are influencing consumers to engage in desired behavior, it is very difficult to determine the ROI of one's advertising spend within and across tactics. Attribution measurement in digital advertising is different from attribution in traditional media, primarily due to consistent user identifiers in the digital ecosystem which allow for user-specific attribution.

1.2.1 User-specific Attribution

Attribution of events to a specific user has been a hallmark of digital advertising, but with varying degrees of complexity and success. The most widely used methodology for single source attribution is the "last click" or "last touch" model, which is considered by many practitioners to represent an incomplete picture of advertising value. It nonetheless has been very widely used due to its simplicity of implementation and the perceived lack of legitimate measurement alternatives. Modern user level attribution systems seek to create a mature valuation of all events in the digital advertising ecosystem by accounting for all the advertising events that a consumer might be exposed to over the course of a campaign.

While technology has allowed for advances in user-specific attribution by enabling the collection and analysis all events across screens (not just the first or last event) it has also introduced new challenges. One key challenge is how to consistently identify the same user across multiple devices as consumers navigate a multi-screen existence. Another key challenge is how to ensure that the rules by which advertising campaigns are optimized—which are now largely executed in real time via automated programmatic platforms—account for the learnings gleaned from user-level cross-platform measurement solutions. Currently, most bidding platforms can only use desktop activity (via cookie-based based measurement) to inform bid multipliers and shift media investments instead of data sets that account for events across all screens like mobile, tablet, and Over-The-Top (OTT) environments.

1.2.2 User-group Attribution

As the digital advertising ecosystem matures, handling campaigns that span both digital and traditional media will require integrating traditional metrics and techniques with digital methods. Often this will mean providing data that is less specific—for example, aggregating click-through rates into demographic or geographic buckets as opposed to providing user-specific ad impressions and clicks. The actual level of aggregation will vary by implementation.

2 Methodologies

2.1 Single Event Attribution

Single source attribution assigns all credit for a desired outcome to a single event. The most basic but most widely used type of attribution is "last touch" attribution, which gives 100% credit to the last meaningful event before a desired outcome takes place, generally the last ad impression (sometimes called ad view), last click, or last engagement.

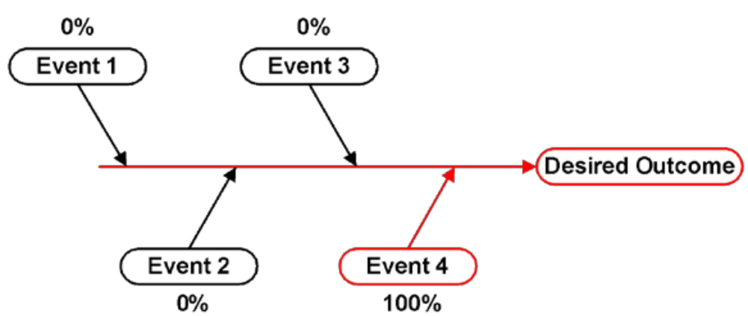

While still widely used, single source attribution has fallen out of favor with many buyers as it does not accurately reflect nor credit all of the contributors to the desired outcome. With this methodology, performance-based-pricing media sellers maximize their revenue by finding prospects furthest down the marketing funnel to "win" the last-touch race.

2.2 Multiple Event Attribution

Multiple source attribution is the process of collecting and analyzing more than one advertising event contributing to an outcome. This type of measurement is based on the belief that all advertising events that occur within a campaign—across channels, platforms, and formats—have a cumulative effect on consumer behavior when contributing to a desired outcome. As such, proponents of multiple source attribution tend to reject the idea that one specific event should claim all the credit for an advertising outcome. Multiple source attribution requires an understanding of the events that occur and of the factors that influence their values.

The set of event types, factors, and models described in this document is not exhaustive, and does not intend to be prescriptive.

2.2.1 Event Types

Clicks
A click is the measurement of navigating from one page to another by activating a hyperlink. Clicks may be broken down into ad clicks, search clicks, affiliate clicks, and possibly other sub-categories.

Engagements

- **Internal Engagements**
  User activity that occurs on the advertiser's directly owned content, but is not the desired outcome. For example, browsing product feature pages without completing a purchase at that time.

- **External Engagements with 3rd party content**
  User activity that occurs on 3rd party owned content, yet promotes the desired outcome in some way. This could include viewing promotional content on an affiliate site, engaging the advertiser's brand on a social media site, or through other methods.

- **External Engagements with 1st party content**
  User activity that occurs in an ad that leverages rich media capabilities. The user interacts with the ad, but this interaction does not result in navigation off the website.

Ad Impression
A measurement of responses from an ad delivery system to an ad request from the user's browser.

Viewable Ad Impression
An ad impression contained in the viewable space of the browser window, on an in-focus browser tab, based on pre-established criteria such as the percent of ad pixels within the viewable space and the length of time the ad is in the viewable space of the browser. It is recognized that an "opportunity to see" the ad exists with a viewable ad impression, which may or may not be the case with a served ad impression.

Direct Navigation
Direct navigation is the act of typing an address into the URL bar directly, or using a bookmark, or otherwise accessing a website without having clicked on a hyperlink.

Search
An event where the user has searched for specific keywords in order to locate something of interest.

2.2.2 Factors

Additionally, each source may be modified by one or more factors:

Recency
How recently an event occurred, whether measured by time or by number of intervening events.

Frequency
How often a specific event occurs.

Sequence position
Whether an event was the first, last, or *n*th in sequence. Determining the "first" event is not perfectly reliable in attribution efforts, as cookie churn and scope issues may mask the true first event.

Path
The serial list of events experienced by a user. See additional information in 2.2.3.

Engagement Depth or Duration
Specific events can indicate different levels of interest by a specific user. When a formal hierarchy is defined for understanding this, the engagement depth can be measured.

2.2.3 Path to Conversion

Once event types are measured along with the additional factors that influence their relative contribution to a desired outcome, the data is generally organized and evaluated in the chronological sequence in which events occurred. This sequence, known as a path to conversion, allows those looking at the information to infer behavioral scenarios that are likely to produce desired outcomes. This analysis can then inform future planning or in-flight optimization between channels, audiences, or creative messaging.

For example, an automobile advertiser trying to increase dealership foot traffic could determine via a path to conversion analysis that more affluent consumers are visiting its locations on weekday evenings after first receiving geo-targeted desktop video ads followed by a mobile display unit showcasing the closest dealership locations for test drives. The automobile advertiser could then surmise that it should increase investment against affluent audiences on desktop video channels on weekday afternoons, and also shift dollars toward mobile display ads on weekday evenings using dealer location creative.

One drawback to path to conversion analysis is that chronological sequences do not provide information about the *degree* to which previous touch points affected the ultimate outcome. As such, advertisers have developed attribution methodologies to weight event types based on assumptions about their relative contribution.

2.2.4 Attribution Methodologies

An attribution methodology is the set of rules by which the value of each event is determined. Here are some example valuation rule sets:

Single Event Attribution – Credit is assigned to a single event along a path to conversion

- **First Touch** – The event receives 100% of the credit if it was the first event recorded. No other events are assigned credit.

- **Last Touch** – The event receives 100% of the credit if it was the last event recorded. No other events are assigned credit.

Fractional - Rules Based – Credit is assigned to multiple events along a path to conversion based on a predetermined set of rules. Examples of rulesets include even weighting, time decay, and u-shaped:

- **Even Weighting** – Credit is applied equally across all events and/or channels measured along a path to conversion. Example: Given ten measureable events along a path to conversion, each is assigned 10% credit.

- **Time Decay** – Credit is applied to events at increasing or decreasing intervals along a path to conversion. Event values are usually altered based on specific time windows when the events occur. Example: 40% of credit could be given to events within 24 hours of conversion, 30% to events within 1-3 days, 20% to events within 3-7 days, and 10% to events within 7-14 days.

- **U-shaped** – Credit is disproportionally applied to events at the beginning and end of a path to conversion. Example: 40% of credit could be given to events occurring in the last day before a desired outcome, 20% to events occurring between days 1-13, and the remaining 40% to events occurring on the first day.

These rules may be determined exclusively by hand, or may adapt automatically over time.

Fractional – Algorithmic – Credit is assigned to multiple events along a path to conversion given computer based, algorithmic analysis of the relationship of events relative to all other events along the path to conversion. Generally speaking, the value calculation of any event can take into account the value of any other event, even if that other event did not lead to the desired outcome. Fractional algorithmic credit is usually determined based on linear regression or game theory concepts.

2.3 Identifying Users Across Screens

Early measurement systems revolved around desktop browser functionality where, at the time, media consumption largely took place on personal or family computers. Cookies were the primary markers to determine when a person was exposed to paid messaging and if a person engaged with the ad unit in a specific way, as well as the events that took place along the path to conversion within a specific campaign.

As media consumption began to fragment across mobile, tablet, and OTT TV platforms, the lack of cookie support within these devices forced industry participants to find new techniques for identifying when the same user sees campaign messaging across different devices and channels. The resulting approach—known as user-level device mapping—attempts to assemble an individual consumer's device graph, largely based on the likelihood that seemingly disparate devices are being used by the same individual. Device graphs are now seen as a necessary foundation for a holistic view of message delivery within a modern, omni-channel digital media campaign.

Device graphs are generally built and maintained by third party analytics organizations that rely on two distinct approaches: probabilistic methods and deterministic methods. Challenges with both approaches are testing accuracy against a consistent baseline and controlling for errors.

2.3.1 Deterministic Approaches

The deterministic method relies on personally identifiable information (PII) to make device matches when a person uses the same persistent identifier—namely an email addresses, a phone number, or credit card information—when logging into an app or website. When a user logs in at any point across multiple devices, deterministic data providers can associate those device IDs in a device graph and use that information to identify or target the same user across multiple screens with great confidence.

Because of the ability to authenticate across devices, deterministic approaches are thought of as the most accurate way to determine user-level device graphs. However, one downside is that this approach cannot control for when other individuals—friends, family, etc.—are using a person's device. Another is the perceived lack of scale across devices, as there are hard limits to the amount of registration data that companies have contingent upon growing user bases.

2.3.2 Probabilistic Approaches

By drawing on aggregation techniques, probabilistic approaches incorporate thousands of anonymous data points—things like device type, operating system, and location data associated with bid requests, time of day, and a host of others—to identify statistically significant correlations between devices. Signals may be also be drawn from known multi-user identifiers like IP addresses or from geographic regions. By using IP to Geo technology—which can establish a ZIP code or other geographical coordinates from an IP address—the incorporation of additional aggregate signals is possible.

Based on these signals, probabilistic techniques attempt to determine the devices that are likely being used by the same person. Once this determination is made, that provider would likely assign a particular statistical ID to the device. For example, if a smart phone, desktop computer and a laptop connect to the same networks or Wi-Fi hotspots at the same time and in the same places every weekday, one can develop a degree of confidence that all three devices belong to a specific person.

Probabilistic approaches are generally considered to be less accurate than deterministic approaches when associating device pairings, as they are largely based on inferred and/or modelled data. One benefit is that these solutions have greater flexibility to scale across devices, meaning that device mappings can potentially incorporate more overall consumer devices than deterministic partners.

2.3.3 Device Mapping Currencies

Device graphs are assembled by associating five primary device currencies: Device IDs, Advertising IDs, Statistical IDs, Cookie IDs, and/or WAN IP addresses. Using publicly available signals, mapping providers need first to be able to consistently identify the same device against these currencies to develop a confidence threshold. The second step is to make an association with other known devices, a process that is often proprietary and used as a primary differentiator by device graphing providers. In addition to device mapping, these currencies can also be used for targeting, segmentation, and/or online-to-offline tracking.

Device IDs, Advertising IDs and Statistical IDs are primarily mobile device markers. The key difference between Device and Advertising IDs is the persistence of the ID with the device. Whereas Device IDs are tied to the hardware or software of the device, Advertising IDs can be reset by consumers so that past behavior is not associated with their new Advertising ID. Advertising IDs are thus considered more user friendly with regard to controlling the use of marketing data. Statistical IDs—determined independently by each individual measurement organization—are inherently different than Device IDs and Advertising IDs in that they are not supplied directly by the device itself. Instead, they are determined and assigned by the measurement provider based on statistical analysis of disparate device signals. This information is generally not consumer facing and thus does not present consumers with an option to "opt-out" of having this information used for marketing purposes.

More information about each currency can be found below:

- **Device ID**: A device-generated identifier set and/or made available by the device's operating system. Users usually cannot control or change a device-generated identifier. Examples include MAC address and UDID.

- **Advertising ID**: A user resettable ID assigned by the device or operating environment for use as an advertising marker. The key difference between a Device ID and an Advertising ID is that Advertising IDs can be reset by the user at any time. Examples include Apple's IDFA and Android's AAID.

- **Statistical ID**: An identifier derived and assigned by an algorithm to determine a device or user based on the values or a combination of standard attributes made available by the device. This analysis is largely dependent upon device information passed in HTTP headers of ad requests, namely device type, operating system, user-agent, fonts, and IP address. Some attributes can change over time due to device changes or updates.

- **WAN IP Address**: While IP addresses are used frequently when developing Statistical IDs, they are also often used as an independent geographic reference point. In particular, a WAN IP—the address used by a household's router—is often used to associate household level TV buy deliveries with a corresponding household level WAN IP address and the devices connected to it.

- **Cookie ID**: A cookie is a small text file and associated alphanumeric identifier generated by a website or a website partner (advertisers, data management platforms, etc). Cookies are stored on a visitor's browser upon arrival at a particular destination, and Cookie IDs are passed along within ad requests. They are most frequently used to determine desktop or laptop associations. Cookies can generally be read only by the assigning service.

2.4 Online to Offline Attribution

Attribution of offline actions (including in-store visits, in-store purchases, and phone calls) to digital actions is an increasing focus of advertisers as new data collection methods become available. Offline purchases still account for the vast majority of consumer activity relative to e-commerce purchases, so understanding online to offline attribution options is of particular interest for advertisers who want to optimize media activity to drive in store sales.

2.4.1 Offline Visitation

One method of online to offline attribution is to associate digital or mobile behavior to in-store foot traffic. Attribution typically occurs by matching the cookie or device identifier associated with an ad exposure on a user's device graph with location data measured through a user's smartphone.

Location data can be gathered from a panel of consumers that have opted in to share their location data with a location-based measurement provider, or through a physical beacon placed in-store.

2.4.2 Offline Purchases

Offline purchases are typically attributed to digital or mobile behavior using a data set, or "audience graph," which associates a cookie or device identifier with an e-mail address or phone number. Advertising exposures are then linked to customer purchases located in an advertiser's CRM.

2.4.3 Phone Calls

Offline purchases are typically attributed to phone calls using a specific call tracking number placed in digital media. Purchases are confirmed by matching call tracking data with customer records, or through speech analytics technology that identifies conversions. View-through attribution can also be done by associating an online identity with a phone call in the same way as the "offline purchase" example.

3 Challenges

3.1 Incomplete Data Set

It is expected that the whole set of user behaviors will not be available in attribution measurement. There are two primary limitations: the coverage of data collection systems, and time.

3.1.1 Limitation Due to Data Collection or Access

There are several circumstances when data sets could be limited for attribution measurement.

- **Walled Gardens** - The development of so-called "walled gardens"—when publishers restrict third party ad technology from operating on their properties—means that third party attribution measurement platforms often cannot collect data about events that happen inside the walled gardens.

- **Incomplete, Limited or Inaccurate Information About User Device Graphs** – User level device graphs are not an exact science, as each vendor has its own methodologies with associated benefits and limitations.

- **Data Access by Individual Media Plan Partners** - The practice of using multiple media partners on a managed-service basis is still relatively common. In this circumstance, optimization of media toward a desired action is usually handled by each partner independent of the others, and often in the absence of a holistic view of overall media activity and the attribution data learnings.

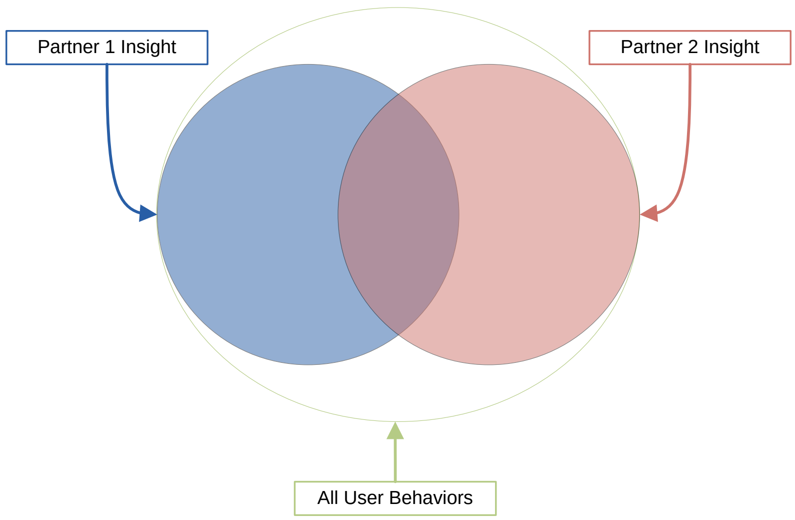

3.1.2 Limitation Due to Time

Storage and processing limits may restrict the look back period for attribution models. Depending on the attribution technology provider, the inclusion of events that occurred more than a certain amount of time before the desired outcome can be limited.

Time also plays a factor in the recognition of the user, because the longer the interval between an event and the desired behavior, the higher the probability of "user identifier churn." "User identifier churn" occurs when existing identifiers, such as cookies, are deleted and replaced by new values.

3.2 User Identifiers

In the short history of digital advertising, the ready availability of user identifiers has often been taken for granted. As the ecosystem has matured, digital advertising has needed to address cross screen scenarios where user identifiers are no longer ubiquitous. This resulted in the probabilistic and deterministic approaches to user level identification outlined above.

The essential areas of growth that best highlight the challenge of user identifiers are the multiple screen scenario, and dealing with online/offline campaigns.

3.2.1 Multiple Screens

As outlined above, a user who accesses the same marketing message across multiple devices will generally have a different user identifier for each device. This means that their engagement with the mobile version of the website cannot be counted in the attribution model for a purchase made via their laptop.

The loss of coverage due to multiple screens is reduced but not entirely mitigated by deterministic and probabilistic solutions to device mapping.

3.2.2 Multiple Media Types

A user who receives a supporting message from traditional media will likely behave differently than a user who does not. However, the traditional media ecosystem does not maintain a unique identifier that can be mapped to the IDs used by digital advertising.

An understanding of the effect of digital advertising in the traditional media ecosystem can be established by mapping user behaviors into compatible buckets (geographic and demographic), and then performing marketing mix modeling.

3.3 Model Scope

Since attribution models are traditionally built using an advertiser's historical advertising data, resulting models tend to be focused on the bottom of the funnel where previous investments were made. Unless the advertiser has been routinely investing across the entire journey—a challenging proposition given the low last-click ROI for upper-funnel channels—it will need to develop a strategy of exploration across the journey to complement its new attribution model. The same approach should also be applied to marketing creatives, ad copy, and website content to ensure that the messages are optimal for the position in the funnel they are reaching.

3.4 Performance Lag

Organizations that engage in daily optimization and reporting may struggle with adopting new attribution models, where investments made upper-funnel may not result in incremental sales until days or weeks later. Marketing leaders will need to build trust in the forward-looking projections of models to defend spend with the promise of future returns. This could be especially problematic during competitive times, such as the retail holiday season where there is less opportunity to correct behavior.

3.5 Managing the Pipeline

While attribution models can reallocate credit to marketing across the funnel, advertisers should use caution when shifting investment at one time. For instance, an advertiser that aggressively shifts credit upper-funnel at the expense of lower-funnel marketing activities may create new customer demand, but then may lack the follow-through investment necessary to drive them the rest of the way through to conversion.

4 Similar Techniques

4.1 Comparison to Marketing Mix Modeling

Like marketing mix modeling, online attribution aims to provide a framework for understanding the valuation of the various consumer-directed messages.

Marketing Mix Modeling	Online Attribution
Aggregate Data Modeling	User-specific Data Modeling, requiring consistent user IDs
Affects Future Spend	Affects Current Spend

4.1.1 Goals of Marketing Mix Modeling

Marketing mix modeling is the use of statistical analysis to optimize future media mix and promotional tactics with respect to sales revenue or profit. It works by modeling large aggregate datasets, and does not require uniquely identified individuals.

4.1.2 Goals of Online Attribution

Online attribution uses the techniques described in this document for several purposes:

- To optimize ongoing campaign delivery based on user-specific data
- To measure the ROI of campaigns
- To determine the amount of payment owed to performance-based pricing vendors

The techniques described are the modeling of multiple signals to understand motivating events leading toward a desired outcome with the specific intent of allocating active revenue based on this understanding.

5 Assessing Results

After a brand selects KPIs and identifies the ideal attribution methodology to measure success, it still remains for advertisers and agencies to analyze the results after the campaign goes into market. This is easier said than done, but nonetheless is always conducted with eye toward identifying useful cross-channel and cross-funnel performance insights to optimize KPI performance over time.

The most common mistake made by campaign strategists is not defining clear success metrics and an attribution strategy *before* launching a campaign. Mid-flight changes to KPIs and measurement approaches can leave big data gaps about what worked and what did not. Similarly, it is important for all parties to understand that different attribution strategies represent different philosophies of understanding campaign success; however, there is no one source of "truth." In a messy real-world setting, no attribution model—not even the most sophisticated ones—can capture a perfect picture of consumer sentiment regarding your campaign. It is a pitfall in any approach to spend enormous amounts of time and resources fixated on the perfect model at the expense of the better.

Lastly, many would argue that there is a degree of fine-tuning of the models themselves that is required to get incrementally more accurate approximations of which campaign elements drive performance over time. A new measurement approach is unlikely to deliver ideal insights out of the box. Rather, tweaking the model using an iterative process, inclusive of many campaigns over a medium-to-long timeframe, is the best strategy to ensure that, whatever attribution methodology or vendor is selected, the results of the model reflect a credible picture of which campaign elements drive business objectives.

6 Validation Methods

When attempting to validate an attribution model, advertisers can pursue a number of different testing methods depending on the complexity of the changes that they are making to their media allocation, as well as the burden of proof required internally to validate the impact of the change—a necessary consideration given that each type of test has its own imperfections.

Geo-testing, for example, is often regarded by our advertisers as the "gold standard" for measuring the impact of an attribution model; however, the structure of the test and high-spend requirements can make it difficult to test more nuanced changes across keywords and campaigns. This often restricts geo-testing to changes in channel-level budgets (search, display, social, etc.) with more granular changes relegated to less precise methods. Marketers also need to be careful of how their bidding platform will respond to the initial geo-segmentation as automatic optimizations can sometimes be affected by the split leading to early biases between the test and control sets.

With lower spend requirements and relatively higher precision, user-based testing is another attractive option. Instead of segmenting by DMA, as in a geo-test, it is possible to expose just a portion of users to incremental spend in a particular area based on the findings of an attribution model. This could be used to validate a particular result of the model, but would be harder to implement holistically.

7 Terminology

- Ad
  A commercial message targeted to an advertiser's customer or prospect.
- Ad Impression
  The delivery of an ad to a web client.
- Click
  The measurement of a navigation from one document to another by activating a hyperlink.
- Direct Navigation
  Direct navigation is the act of typing in a URL into the address bar of a web browser.
- Engagement
  Content-specific metric that evaluates the interest level of a user.

- Engagement Depth
  Specific events can indicate different levels of interest by a specific user. When a formal hierarchy is defined for understanding this, the engagement depth is measured.
- Event
  Any user action that is measured.
- Factor
  A supplementary data point that modifies the impact of the event with which it is associated.
- Frequency
  How often a specific event occurs.
- IP Address
  Generally, the IPv4 address of an internet-connected device.
- IP to Geo
  The process of establishing a ZIP code or other geographic coordinate from an IP address.
- Path
  The serial list of events experienced by a user.
- Recency
  How recently an event occurred, whether measured by time or by number of intervening events.
- Scope
  The set of data available to a participant.
- Search
  An event where the user has provided keywords to a form in order to locate something of interest.
- Sequence position
  Whether an event was the first, last, or n^{th} in sequence. Determining the "first" event is not perfectly reliable in attribution efforts, since cookie churn and scope issues may mask the true first event.