



- (51) **International Patent Classification:**
G06Q 30/02 (2012.01)
- (21) **International Application Number:**
PCT/US2013/023186
- (22) **International Filing Date:**
25 January 2013 (25.01.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/363,913 1 February 2012 (01.02.2012) US
- (71) **Applicant:** GOOGLE INC. [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).
- (72) **Inventors:** HUGHES, Gabriel; 6 Lynmouth Road, London E17 8AG (GB). KENT, James; Flat 16, Goddard House, 3 George Mathers Road, London SE11 4BG (GB).
- (74) **Agents:** LANZA, John D. et al.; FOLEY & LARDNER LLP, 3000 K Street N.W., Suite 600, Washington, District of Columbia 20007-5109 (US).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

[Continued on next page]

(54) Title: ALTERATIONS OF CALCULATIONS IN ATTRIBUTION MODELING

(57) Abstract: A computer system provides attribution based on advertisement conversion data includes a processing circuit. The processing circuit is configured to receive user interaction data; determine that a conversion event has occurred based on user interaction data and conversion criteria; store conversion path data based on user interaction data; attribute the conversion event at least in part to a category in the conversion path using an attribution model to generate attribution data; and generate a marketing purchase parameter based on output of the attribution model.




FIG. 2



Published:

— *with international search report (Art. 21(3))*

ALTERATIONS OF CALCULATIONS IN ATTRIBUTION MODELING

BACKGROUND

[0001] The present disclosure relates generally to the field of advertisements. More specifically, the present disclosure relates to attribution modeling in advertising.

[0002] The internet provides access to a wide variety of content. For instance, images, audio, video, and webpages for a myriad of different topics are accessible through the Internet. The accessible content provides an opportunity to place advertisements. Advertisements can be placed within content, such as a webpage, image or video, or the content can trigger the display of one or more advertisements, such as presenting an advertisement in an advertisement slot.

[0003] Advertisers decide which ads are displayed within particular content using various advertising management tools. These tools also allow an advertiser to track the performance of various ads or ad campaigns. The parameters used to determine when to display a particular ad can also be changed using advertising management tools. A user often is exposed to or interacts with more than one marketing category prior to a conversion event. The categories may relate to a grouping of marketing activity, such as a grouping based on a marketing channel, or a grouping based on a type of keyword used (e.g., a brand term versus a generic term).

[0004] Attribution modeling is the practice of attributing credit to marketing categories that led to a website and subsequently resulted in a conversion event. An attribution model includes the algorithm which determines how conversion credit is shared between multiple marketing categories.

SUMMARY

[0005] One implementation of the present disclosure relates to a computer system for providing attribution based on advertisement conversion data. The system includes a processing circuit configured to receive user interaction data, wherein the user interaction data specifies user interaction with content items and conversion items. The processing circuit is further configured to determine that a conversion event has occurred based on the user interaction data and the conversion criteria. The processing circuit is further configured to store conversion path data based on the user interaction data.. The processing circuit is

further configured to attribute the conversion event at least in part to a category in the conversion path using an attribution model to generate attribution data. The processing circuit is further configured to generate a marketing purchase parameter based on output of the attribution model.

[0006] Another implementation of the present disclosure relates to a method for providing attribution based on advertisement conversion data. The method includes receiving user interaction data. The method further includes determining that a conversion event has occurred based on the user interaction data and the conversion criteria. The method further includes storing conversion path data based on the user interaction data. The method further includes attributing the conversion event at least in part to a category in the conversion path using an attribution model to generate attribution data. The method further includes generating a marketing purchase parameter based on output of the attribution model.

[0007] Another implementation of the present disclosure relates to one or more computer readable media having instructions stored therein, the instructions being executable by one or more processors to cause the one or more processors to perform operations. The instructions include receiving user interaction data. The instructions further include determining that a conversion event has occurred based on the user interaction data and the conversion criteria. The instructions further include storing conversion path data based on the user interaction data. The instructions further attributing the conversion event at least in part to a category in the conversion path using an attribution model to generate attribution data. The instructions further include generating a marketing purchase parameter based on output of the attribution model.

[0008] These implementations are mentioned not to limit or define the scope of the disclosure, but to provide an example of an implementation of the disclosure to aid in understanding thereof. Particular implementations may be developed to realize one or more of the following advantages.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The details of one or more implementations of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

[0010] FIG. 1 is a block diagram of an example environment in which an advertisement management system manages advertising services in accordance with a described implementation.

[0011] FIG. 2 is a flow diagram of a process for using multiple attribution models to report return on advertising spend in accordance with a described implementation.

[0012] FIG. 3 is an illustration of conversion path data in accordance with a described implementation.

[0013] FIG. 4 is an illustration of a last click attribution model in accordance with a described implementation.

[0014] FIG. 5 is an illustration of a first click attribution model in accordance with a described implementation.

[0015] FIG. 6 is an illustration of a first touch attribution model in accordance with a described implementation.

[0016] FIG. 7 is an illustration of a category touch attribution model in accordance with a described implementation.

[0017] FIG. 8 is an illustration of a recency attribution model in accordance with a described implementation.

[0018] FIG. 9 is an illustration of exemplary display data or report data showing results of multiple attribution model calculations in accordance with a described implementation.

[0019] FIG. 10A is an illustration of exemplary display data or report data showing results of multiple attribution model calculations comprising cost data in accordance with a described implementation.

[0020] FIG. 10B is an illustration of exemplary display data or report data showing results of multiple attribution model calculations comprising cost data in accordance with a described implementation.

[0021] FIG. 11 is a block diagram of an exemplary computer system for use with the described implementation.

[0022] FIG. 12 is a flow diagram of a process for providing attribution based on advertisement conversion data in accordance with a described implementation.

[0023] FIG. 13 a flow diagram of a process 1200 for providing attribution based on advertisement conversion data in accordance with a described implementation.

[0024] FIGS. 14-15 are illustrations of an exemplary process of re-scaling and re-weighting an attribution model in accordance with a described implementation.

DETAILED DESCRIPTION OF THE EXEMPLARY EMBODIMENTS

[0025] In a described implementation, a plurality of different attribution models or algorithms may be applied to conversion path data. Each model may provide differing credit to the different categories in the conversion path. In one or more implementations, cost data for each category may be processed with the differing credit values to provide information about the relative values of using each category in a marketing campaign. Advertisers or other content providers may then use this value information to improve their marketing campaigns by, e.g., allocating more or less money to the different categories in a future campaign.

[0026] Advertisers and content providers may be provided various reports that disclose various user interactions with ads and content. Each user interaction can include a number of dimensions, which can contain data associated with the user interaction. Reports can be generated to provide an advertiser with information regarding the user interactions. User interactions can include user interactions from various categories. The categories may relate to a grouping of marketing activity, such as a grouping based on a marketing channel, or a grouping based on a type of keyword used (e.g., a brand term versus a generic term). Channels are a way to describe the originating source of a user interaction. Illustrative examples of user interactions and channels include clicking on a paid advertisement, directly navigating to a website, clicking on an organic search result, clicking on a link within an email, clicking on a link from a referring website, clicking on a link from a social networking website, mousing over an advertisement, providing a display on a webpage of a banner advertisement or other advertisement to give the user an opportunity to see the advertisement in a case where the user does not click on it or otherwise positively interact with it, etc. Conversion paths include one or more user interactions that preceded a conversion user interaction.

[0027] User interactions include any presentation of content to a user and any subsequent affirmative actions or non-actions (collectively referred to as “actions” unless otherwise specified) that a user takes in response to presentation of content to the user (e.g., selections of the content following presentation of the content, or no selections of the content following the presentation of the content). Thus, a user interaction does not necessarily require a selection of the content (or any other affirmative action) by the user. A user interaction may be merely an exposure or impression that a user views or does not actually view. User interaction data may include user exposures, for example to include data regarding a user

who is tracked as being exposed to an ad impression but does not click or proactively interact with it.

[0028] User interaction measures can include one or more of time lag measures (i.e., measures of time from one or more specified user interactions to a conversion), path length measures (i.e., quantities of user interactions that occurred prior to conversions), user interaction paths (i.e., sequences of user interactions that occurred prior to the conversion), assist interaction measures (i.e., quantities of particular user interactions that occurred prior to the conversion), and assisted conversion measures (i.e., quantities of conversions that were assisted by specified content). User interaction data may be collected from serving and tracking activity across the Internet and across many networks. Tracking tags may be used. A user may opt out of this tracking activity in a number of ways. Data may be collected for a plurality of different attribution models.

[0029] Attribution modeling

[0030] FIG. 1 is a block diagram of an example environment 100 in which an advertisement management system manages advertising services, in accordance with a described implementation. Example environment 100 includes a network 102, such as a local area network (LAN), a wide area network (WAN), the Internet, or a combination thereof. Network 102 connects websites 104, user devices 106, advertisers 108, and an advertisement management system 110. Environment 100 may include many thousands or more of websites 104, user devices 106 and advertisers 108.

[0031] A website 104 includes one or more resources 105 associated with a domain name and hosted by one or more servers. An example website is a collection of webpages formatted in hypertext markup language (HTML) that can contain text, images, multimedia content, and programming elements, such as scripts.

[0032] When a user is exposed to or interacts with a marketing channel, the channel has a number of attributes that may be of interest to marketing analysis, such as the site, the type of ad, the ad campaign which the ad forms a part of, the site placement (which page on a site), the particular ad message, the particular ad image or text, and the ad size. When analyzing the value of advertising using attribution models, the differences in these attributes can be compared through crediting all like attributes using an attribution model. Thus, a given attribution model may be used to compare the attributed credit for all sites included in the analysis; but may also be used to compare all campaigns, or to compare all ad placements, and so on, depending on the attribute which is being analyzed.

[0033] A resource 105 is any data that can be provided over network 102. Resource 105 may be identified by a resource address that is associated with the resource 105, such as a uniform resource locator (URL). Resources 105 can include webpages, word processing documents, portable document format (PDF) documents, images, video, programming elements, interactive content, and feed sources, to name only a few. Resources 105 can include content, such as words, phrases, images and sounds, that may include embedded information (such as meta-information in hyperlinks) and/or embedded instructions.

Embedded instructions can include code that is executed at a user's device, such as in a web browser. Code can be written in languages such as JavaScript® or ECMAScript®.

[0034] User device 106 is an electronic device that is under control of a user and is capable of requesting and receiving resources 105 over network 102. Example user devices 106 include personal computers, mobile communication devices, and other devices that can send and receive data over network 102. User device 106 typically includes a user application, such as a web browser, to facilitate the sending and receiving of data over network 102.

[0035] User device 106 can request resources 105 from a website 104. In turn, data representing resource 105 can be provided to user device 106 for presentation by user device 106. The data representing resource 105 can include data specifying a portion of the resource or a portion of a user display (e.g., a presentation location of a pop-up window or in a slot of a webpage) in which advertisements can be presented. These specified portions of resource 105 or user display are referred to as advertisement slots.

[0036] To facilitate searching of the vast number of resources 105 accessible over the network 102, environment 100 can include a search system 112 that identifies resources 105 by crawling and indexing resources 105 provided on websites 104. Data about resources 105 can be indexed based on resource 105 with which the data is associated. The indexed and, optionally, cached copies of resources 105 are stored in a search index (not shown).

[0037] User devices 106 can submit search queries to search system 112 over network 102. In response, search system 112 accesses the search index to identify resources 105 that are relevant to the search query. In one implementation, a search query includes one or more keywords. Search system 112 identifies resources 105 that are responsive to the query, provides information about resources 105 in the form of search results and returns the search results to user devices 106 in search results pages. A search result can include data generated by search system 112 that identifies a resource 105 that is responsive to a particular search query, and can include a link to resource 105. An example search result can include a webpage title, a snippet of text or a portion of an image extracted from webpage 104, a

rendering of resource 105, and the URL of webpage 104. Search results pages can also include one or more advertisement slots in which advertisements can be presented.

[0038] A search result page can be sent with a request from search system 112 for the web browser of user device 106 to set an HTTP (HyperText Transfer Protocol) cookie. A cookie can represent, for example, a particular user device 106 and a particular web browser. For example, search system 112 includes a server that replies to the query by sending the search results page in an HTTP response. This HTTP response includes instructions (e.g., a set cookie instruction) that cause the browser to store a cookie for the site hosted by the server or for the domain of the server. If the browser supports cookies and cookies are enabled, every subsequent page request to the same server or a server within the domain of the server will include the cookie. The cookie can store a variety of data, including a unique or semi-unique identifier. The unique or semi-unique identifier can be anonymized and is not connected with user names. Because HTTP is a stateless protocol, the use of cookies allows an external service, such as search system 112 or other system, to track particular actions and status of a user over multiple sessions. A user may opt out of tracking user actions, for example, by disabling cookies in the browser's settings.

[0039] When a resource 105 or search results are requested by user device 106 or provided to user device 106, advertisement management system 110 receives a request for advertisements to be provided with resource 105 or search results. The request for advertisements can include characteristics of the advertisement slots that are defined for the requested resource 105 or search results page, and can be provided to advertisement management system 110. For example, a reference (e.g., URL) to resource 105 for which the advertisement slot is defined, a size of the advertisement slot, and/or media types that are available for presentation in the advertisement slot can be provided to advertisement management system 110. Similarly, keywords (i.e., one or more words that are associated with content) associated with a requested resource 105 ("resource keywords") or a search query for which search results are requested can also be provided to advertisement management system 110 to facilitate identification of advertisements that are relevant to resource 105 or a search query.

[0040] Based on data included in the request for advertisements, advertisement management system 110 can select advertisements that are eligible to be provided in response to the request ("eligible advertisements"). For example, eligible advertisements can include advertisements having characteristics matching the characteristics of advertisement slots and that are identified as relevant to specified resource keywords or search queries. In some

implementations, advertisements having targeting keywords that match the resource keywords, the search query, or portions of the search query are selected as eligible advertisements by advertisement management system 110.

[0041] Various mechanisms may be used to select ads for presentation. According to an example embodiment, relevant advertisements submitted by advertisers may be considered for ad auctions that are performed whenever a user performs an online search using an online search application or whenever a user visits a content page with available advertisement slots. The keywords associated with the user's online search or with the content page are compared to keywords and phrases associated with ads of an advertiser (as entered by the advertiser), and if there is a sufficient match, and based on other parameters such as predicted click thru rate (pCTR) of the ad and the bid entered for that ad by the advertiser, the ad may either win the auction or lose the auction to other competing ads. If the ad wins the auction, that ad is displayed to the user along with the search results or content page. For the user actions that cause ad auctions to be performed, a combination of bid and quality is used by an ad server to determine which ads in a database of candidate ads "win" each auction and are thus shown to the user, whereby all ads that "lose" each auction are returned by the ad server back to the ad indexes that store the ads for consideration in current and future ad auctions.

[0042] The advertisement management system may include an automated bidding system made available to advertisers that are configuring ad campaigns. The automated bidding system may be a tool configured to generate marketing purchase parameters for an advertiser based on parameters set by the advertiser. For purposes of providing an example, the discussion below describes an arrangement in which the marketing purchase parameter is a keyword bid. It will be appreciated, however, that other marketing purchase parameters may be used. For example, the marketing purchasing parameter may also be a maximum cost, a target cost (e.g., a target cost per action, a target cost per mille, a target cost per click, etc.), and so on. As other examples, the marketing purchasing parameter may relate to a bid for placement/display of an advertisement in a specified manner. For example, the advertiser may specify that an ad should be shown ten times per day. The automated bidding system may then increase or decrease the bids for the ad to achieve ten impressions per day. Budget limits may also be set within the automated bidding system (e.g., to ensure that the total amount spent on advertising does not exceed a specified threshold).

[0043] Advertisement management system 110 selects an eligible advertisement for each advertisement slot of resource 105 or of a search results page. Resource 105 or search results page is received by user device 106 for presentation by user device 106. User interaction

data representing user interactions with presented advertisements can be stored in a historical data store 119. For example, when an advertisement is presented to the user via an ad server 114, data can be stored in a log file 116. This log file 116, as more fully described below, can be aggregated with other data in historical data store 119. Accordingly, historical data store 119 contains data representing the advertisement impression. For example, the presentation of an advertisement is stored in response to a request for the advertisement that is presented. For example, the ad request can include data identifying a particular cookie, such that data identifying the cookie can be stored in association with data that identifies the advertisement(s) that were presented in response to the request. In some implementations, the data can be stored directly to historical data store 119.

[0044] Similarly, when a user selects (i.e., clicks) a presented advertisement, data representing the selection of the advertisement can be stored in log file 116, a cookie, or historical data store 119. In some implementations, the data is stored in response to a request for a webpage that is linked to by the advertisement. For example, the user selection of the advertisement can initiate a request for presentation of a webpage that is provided by (or for) the advertiser. The request can include data identifying the particular cookie for the user device, and this data can be stored in the advertisement data store.

[0045] User interaction data can be associated with unique identifiers that represent a corresponding user device with which the user interactions were performed. For example, in some implementations, user interaction data can be associated with one or more cookies. Each cookie can include content which specifies an initialization time that indicates a time at which the cookie was initially set on the particular user device 106.

[0046] Log files 116, or historical data store 119, also store references to advertisements and data representing conditions under which each advertisement was selected for presentation to a user. For example, historical data store 119 can store targeting keywords, bids, and other criteria with which eligible advertisements are selected for presentation. Additionally, historical data store 119 can include data that specifies a number of impressions for each advertisement and the number of impressions for each advertisement can be tracked, for example, using the keywords that caused the advertisement impressions and/or the cookies that are associated with the impressions. Data for each impression can also be stored so that each impression and user selection can be associated with (i.e., stored with references to and/or indexed according to) the advertisement that was selected and/or the targeting keyword that caused the advertisement to be selected for presentation.

[0047] Advertisers 108 can submit, to advertisement management system 110, campaign parameters (e.g., targeting keywords and corresponding bids) that are used to control distribution of advertisements. Advertisers 108 can access advertisement management system 110 to monitor performance of the advertisements that are distributed using the campaign parameters. For example, an advertiser can access a campaign performance report that provides a number of impressions (i.e., presentations), selections (i.e., clicks), and conversions that have been identified for the advertisements. The campaign performance report can also provide a total cost, a cost-per-click, and other cost measures for the advertisement over a specified period of time. For example, an advertiser may access a performance report that specifies that advertisements distributed using the phrase match keyword “hockey” have received 1,000 impressions (i.e., have been presented 1,000 times), have been selected (e.g., clicked) 20 times, and have been credited with 5 conversions. Thus, the phrase match keyword hockey can be attributed with 1,000 impressions, 20 clicks, and 5 conversions.

[0048] As described above, reports that are provided to a particular content provider can specify performance measures measuring user interactions with content that occur prior to a conversion. A conversion occurs when a user performs a specified action, and a conversion path includes a conversion and a set of user interactions occurring prior to the conversion by the user. What constitutes a conversion may vary from case to case and can be determined in a variety of ways. For example, a conversion may occur when a user clicks on an advertisement, is referred to a webpage or website, and then consummates a purchase there before leaving the webpage or website. As another example, a conversion may occur when a user spends more than a given amount of time on a particular website. Data from multiple user interactions can be used to determine the amount of time at the particular website.

[0049] Actions that constitute a conversion can be specified by each advertiser. For example, each advertiser can select, as a conversion, one or more measurable/observable user actions such as, for example, downloading a white paper, navigating to at least a given depth of a website, viewing at least a certain number of webpages, spending at least a predetermined amount of time on a website or webpage, purchasing a product, or registering on a website. Other actions that constitute a conversion can also be used.

[0050] To track conversions (and other interactions with an advertiser’s website), an advertiser can include, in the advertiser’s webpages, embedded instructions that monitor user interactions (e.g., page selections, content item selections, and other interactions) with advertiser’s website, and can detect a user interaction (or series of user interactions) that

constitutes a conversion. In some implementations, when a user accesses a webpage, or another resource, from a referring webpage (or other resource), the referring webpage (or other resource) for that interaction can be identified, for example, by execution of a snippet of code that is referenced by the webpage that is being accessed and/or based on a URL that is used to access the webpage.

[0051] For example, a user can access an advertiser's website by selecting a link presented on a webpage, for example, as part of a promotional offer by an affiliate of the advertiser. This link can be associated with a URL that includes data (i.e., text) that uniquely identifies the resource from which the user is navigating. For example, the link http://www.example.com/homepage/%affiliate_identifier%promotion_1 specifies that the user navigated to the example.com webpage from a webpage of the affiliate that is associated with the affiliate identifier number that is specified in the URL, and that the user was directed to the example.com webpage based on a selection of the link that is included in the promotional offer that is associated with promotion_1. The user interaction data for this interaction (i.e., the selection of the link) can be stored in a database and used, as described below, to facilitate performance reporting.

[0052] When a conversion is detected for an advertiser, conversion data representing the conversion can be transmitted to a data processing apparatus ("analytics apparatus") that receives the conversion data, and in turn, stores the conversion data in a data store. This conversion data can be stored in association with one or more cookies for the user device that was used to perform the user interaction, such that user interaction data associated with the cookies can be associated with the conversion and used to generate a performance report for the conversion.

[0053] Typically, a conversion is attributed to a targeting keyword when an advertisement that is targeted using the targeted keyword is the last clicked advertisement prior to the conversion. For example, advertiser X may associate the keywords "tennis," "shoes," and "Brand-X" with advertisements. In this example, assume that a user submits a first search query for "tennis," the user is presented a search result page that includes advertiser X's advertisement, and the user selects the advertisement, but the user does not take an action that constitutes a conversion. Assume further that the user subsequently submits a second search query for "Brand-X," is presented with the advertiser X's advertisement, the user selects advertiser X's advertisement, and the user takes action that constitutes a conversion (e.g., the user purchases Brand-X tennis shoes). In this example, the keyword "Brand-X" will be credited with the conversion because the last advertisement selected prior to the conversion

(“last selected advertisement”) was an advertisement that was presented in response to the “Brand-X” being matched.

[0054] Providing conversion credit to the keyword that caused presentation of the last selected advertisement (“last selection credit”) prior to a conversion is a useful measure of advertisement performance, but this measure alone does not provide advertisers with data that facilitates analysis of a conversion cycle that includes user exposure to, and/or selection of, advertisements prior to the last selected advertisement. For example, last selection credit measures alone do not specify keywords that may have increased brand or product awareness through presentation of advertisements that were presented to, and/or selected by, users prior to selection of the last selected advertisement. However, these advertisements may have contributed significantly to the user subsequently taking action that constituted a conversion.

[0055] In the example above, the keyword “tennis” is not provided any credit for the conversion, even though the advertisement that was presented in response to a search query matching the keyword “tennis” may have contributed to the user taking an action that constituted a conversion (e.g., making a purchase of Brand-X tennis shoes). For instance, upon user selection of the advertisement that was presented in response to the keyword “tennis” being matched, the user may have viewed Brand-X tennis shoes that were available from advertiser X. Based on the user’s exposure to the Brand-X tennis shoes, the user may have subsequently submitted the search query “Brand-X” to find the tennis shoes from Brand-X. Similarly, the user’s exposure to the advertisement that was targeted using the keyword “tennis,” irrespective of the user’s selection of the advertisement, may have also contributed to the user subsequently taking action that constituted a conversion (e.g., purchasing a product from advertiser X). Analysis of user interactions, with an advertiser’s advertisements (or other content), that occur prior to selection of the last selected advertisement can enhance an advertiser’s ability to understand the advertiser’s conversion cycle.

[0056] A conversion cycle is a period that begins when a user is presented an advertisement and ends at a time at which the user takes action that constitutes a conversion. A conversion cycle can be measured and/or constrained by time or actions and can span multiple user sessions. User sessions are sets of user interactions that are grouped together for analysis. Each user session includes data representing user interactions that were performed by a particular user and within a session window (i.e., a specified period). The session window can be, for example, a specified period of time (e.g., 1 hour, 1 day, or 1 month) or can be delineated using specified actions. For example, a user search session can include user search

queries and subsequent actions that occur over a 1 hour period and/or occur prior to a session ending event (e.g., closing of a search browser).

[0057] Analysis of a conversion cycle can enhance an advertiser's ability to understand how its customers interact with advertisements over a conversion cycle. For example, if an advertiser determines that, on average, an amount of time from a user's first exposure to an advertisement to a conversion is 20 days, the advertiser can use this data to infer an amount of time that users spend researching alternative sources prior to converting (i.e., taking actions that constitute a conversion). Similarly, if an advertiser determines that many of the users that convert do so after presentation of advertisements that are targeted using a particular keyword, the advertiser may want to increase the amount of money that it spends on advertisements distributed using that keyword and/or increase the quality of advertisements that are targeted using that particular keyword.

[0058] Measures of user interactions that facilitate analysis of a conversion cycle are referred to as conversion path performance measures. Conversion path performance measures specify durations of conversion cycles, numbers of user interactions that occurred during conversion cycles, paths of user interactions that preceded a conversion, numbers of particular user interactions that occurred preceding conversions, as well as other measures of user interaction that occurred during conversion cycles, as described in more detail below.

[0059] Advertisement management system 110 includes a performance analysis apparatus 120 that determines conversion path performance measures that specify measures of user interactions with content items during conversion cycles. Performance analysis apparatus 120 tracks, for each advertiser, user interactions with advertisements that are provided by the advertiser, determines (i.e., computes) one or more conversion path performance measures, and provides data that cause presentation of a performance report specifying at least one of the conversion path performance measures. Using the performance report, the advertiser can analyze its conversion cycle, and learn how each of its keywords cause presentation of advertisements that facilitate conversions, irrespective of whether the keywords caused presentation of the last selected advertisement. In turn, the advertiser can adjust campaign parameters that control distribution of its advertisements based on the performance report.

[0060] Configuration options can be offered to reduce bias in performance reports. Without configuration options, some performance reports can be biased, such as towards short conversion paths. For example, a performance report can be biased towards short conversion paths if data used as a basis for the report includes a percentage of partial conversion paths which is higher than a threshold percentage. A partial conversion path is a

conversion path in which some but not all user interaction data for a user is associated with a conversion. A partial conversion path can be included in a report if, for example, the report is generated using a reporting period which is less than the length of a typical conversion cycle for the advertiser who requested the report.

[0061] A reporting period determines the maximum length (in days) of a reported conversion cycle because additional data outside of the reporting period is not used to generate the report. A performance report can be based on a reporting period (i.e., lookback window), such that user interactions prior to the reporting period are not considered part of the conversion cycle when generating the report. Such a reporting period is referred to as a “lookback window”. For example, when generating a report with a lookback window of thirty days, available user interaction data representing user actions that occurred between July 1 and July 31 of a given year would be available for a conversion that occurred on July 31 of that year.

[0062] If a default lookback window (e.g., thirty days) is used, the performance report can be biased towards short conversion paths if the typical conversion cycle length for a product associated with the report is greater than the default lookback window. For instance, in the example above, a typical conversion cycle for "Brand-X" tennis shoes may be relatively short (e.g., thirty days) as compared to a conversion cycle for a more expensive product, such as a new car. A new car may have a much longer conversion cycle (e.g., ninety days).

[0063] An attribution model may include an algorithm defining which of a plurality of lookback windows to use and/or modifying these lookback windows based on user ad exposure, interaction and conversion data. In one example, different attribution models may apply different lookback windows to determine conversion events. The use of lookback windows to select conversion events may be kept as broad as possible such that all conversion events in a given time period are selected for processing by the attribution model algorithms.

[0064] Different advertisers or different products for an advertiser can have different associated conversion cycle lengths. For example, an advertiser that sells low cost (e.g., less than \$100) products may specify a lookback window of 30 days, while an advertiser that sells more expensive products (e.g., at least \$1000) may specify a lookback window of 90 days.

[0065] In some implementations, advertiser 108 can specify a lookback window to use when requesting a performance report, such as by entering a number of days or by selecting a lookback window from a list of specific lookback windows (e.g., thirty days, sixty days, ninety days). Allowing an advertiser to configure the lookback window of their performance

reports enables the advertiser to choose a lookback window that corresponds to conversion cycles of their products. Allowing lookback window configuration also enables advertisers to experiment with different lookback windows, which can result in the discovery of ways to improve conversion rates.

[0066] Other factors can contribute to reporting on partial conversion paths. For example, as mentioned above, user interaction data used as a basis for a report can be associated with unique identifiers that represent a user device with which the user interactions were performed. As described above, a unique identifier can be stored as a cookie. Cookies can be deleted from user devices, such as by a user deleting cookies, a browser deleting cookies (e.g., upon browser exit, based on a browser preference setting), or some other software (e.g., anti-spyware software) deleting cookies.

[0067] If cookies are deleted from a user device, a new cookie will be set on the user's device when the user visits a webpage (e.g., search system 112). The new cookie may be used to store a new quasi-unique identifier, and thus subsequent user interaction data that occurs on the user device may be associated with a different identifier. Therefore, because each user identifier is considered to represent a different user, the user interaction data associated with the deleted cookies are identified as being associated with a different user than the user interaction data that is associated with the new cookies.

[0068] For instance, in the example above, assume that the user deletes cookies after the first search query for "tennis" is performed and that the second search query for "Brand-X" occurs after the cookies are deleted. In this example, performance measures computed based on the user interaction data for the user can show a bias. For example, a path length measure can be computed as one, rather than two, since the advertisement selection resulting from the first search query is not considered part of the same conversion cycle as the advertisement selection resulting from the second search query, since the two user interactions do not appear to have been performed by the same user.

[0069] To view a report which reduces bias caused from partial conversion paths, an advertiser can specify a lookback window for the report. As described above, the lookback window specifies that the user interaction data used to generate the report are user interaction data that are associated with unique identifiers that have initialization times that are prior to a specified period (e.g., thirty days, sixty days, ninety days) before the conversions. Thus, conversions for which user interaction data that are associated with unique identifiers having initialization times that are after the specified period are excluded from inclusion as a basis for the report. A unique identifier that has a recent initialization time indicates that the

unique identifier may have been recently reinitialized on the user device that the unique identifier represents. Accordingly, user interaction data associated with the relatively new unique identifier may represent only a partial conversion path. Alternatively, conversions for which user interaction data that are associated with unique identifiers having initialization times that are after the specified period are included in the report. To reduce bias, any user interaction included in the conversion path that occurred after the specified period are removed from the conversion path prior to being included in the report.

[0070] Advertisement management system 110, performance analysis apparatus 120, or another computing device may be configured to perform a process for integrating user interaction log data. The process updates conversion paths and determines conversions based upon the updated conversion paths of users.

[0071] As described above, log files 116 may contain user interaction data. A log file 116 may be combined with user interaction data from other logs from other servers, including those that implement search system 112, prior to processing. It may be determined that a new log is available for processing. For example, a notification can be sent to the computing device indicating that a new log is ready for processing, or the existence of a new log can indicate that the new log is ready for processing.

[0072] Next, the new log is retrieved. The new log may be retrieved over network 102. The stateful history for each user is updated based upon the user activity indicated by the new log. The new log can contain information relating to user interactions for numerous users. Historical data store 119 contains user interaction data from previously processed log files. The user interaction data contained within historical data store 119 can be stateful, in that the user interaction data can be grouped by user identifier and ordered chronologically.

[0073] The new log can contain user interaction data for one or more user identifiers. The user interaction data can be grouped by user identifiers and then sorted chronologically. The grouped and sorted user interaction data can then merged with the user interaction data stored in historical data store 119. If a user identifier previously existed in historical data store 119, the new user interaction data are added to the previous user interaction data. Otherwise, the new user interaction data is added with a new user identifier.

[0074] Based upon the updated user interaction data, any conversions that occurred in each of the updated paths of user interactions can be determined. User interaction paths are constrained to those user interactions that are related to a particular advertiser 108. The conversion interactions of the particular advertiser 108 are used to determine if a conversion

has occurred. Once found, the conversion paths can be written to another portion of historical data store 119 or another data store for further analysis.

[0075] Each user interaction includes a set of data or dimensions associated with the user interaction. The dimensions can be sparsely populated, such that, any user interaction may have data relating to a subset of the dimensions. A large number of conversion paths can be generated based upon received user interaction data. Various reports regarding how a campaign or an advertiser's placements are performing can include various information regarding the conversion paths. Given the large potential number of conversion paths, various conversion paths can be grouped together to reduce the number of distinct conversion paths that are reported. In an implementation, conversion paths that have the same number of user interactions and have corresponding data can be aggregated.

[0076] In one implementation, users are able to create various groups to classify individual user interactions. A group includes a group definition that includes one or more group rules that determine if a particular user interaction belongs to a particular group. The group rules use the dimensional data of the user interaction to determine the group of a user interaction. Boolean operators such as AND, NOT, OR, etc. can be used to join various group rules in a group definition. Each group also includes a group name. In some implementations, a group can include display information, such as, but not limited to, a text color and/or background color used to display the group name. Default groups may also be available to users. When default groups are available, a user can copy a default group, including the associated group rules, and then modify one or more of the group rules and/or the group name. User created groups can be stored in a data store, such as a local or remote database. The groups can then be accessed, modified, or deleted at a later time.

[0077] One or more groups can be associated with one another in a sorted or ordered list of grouping definitions. The groups within the ordered list are used to determine the group for each user interaction. The ordering of the list determines the priority of a particular group. A user interaction is grouped with the matching group that has the highest priority. A matching group of a lower priority will be ignored.

[0078] Using the ordered list of grouping definitions, each conversion path can be converted into a group path. A group path contains group elements that correspond to the user interactions of a conversion path. The group element can contain or reference data from the corresponding user interaction. In addition, the group element contains or references the group name and display information of the matching group.

[0079] In one implementation, conversion paths are converted into group paths by adding a reference to the matching group to each of the user interactions. In another implementation, group paths that are separate from the conversion paths are created. In this implementation, the group paths can be stored in the same or in a different location from the location where the conversion paths are stored. Regardless of how the group paths are implemented, the group paths can be aggregated based upon the length of the group path and the group name of the group elements that make up the group path.

[0080] In one implementation, the group paths contain various data from the corresponding conversion path. For example, a conversion path can include a monetary value associated with the conversion. As the group paths are aggregated, the value of all conversion paths associated with the aggregated group paths can also be aggregated. This aggregated value can be included in a report.

[0081] Referring now to FIG. 2, a flow chart illustrating a method of using multiple attribution models to report return on advertising spend will be described. The algorithm which implements the steps of FIG. 2 may be operable on one or more of advertisement management system 110 (FIG. 1), performance analysis apparatus 120, or other components shown in FIG. 1. At block 202, the system is configured to receive user interaction data. At block 204, conversion criteria are stored and access to the conversion criteria is provided to the system. At block 206, the system is configured to determine whether a conversion event has occurred based on the user interaction data and the conversion criteria. Regardless of whether a conversion event has occurred, partial or complete conversion path data is stored at block 208.

[0082] At block 210, the system is configured to provide attribution of one or more conversion events detected in block 206. Attribution may be provided to one or more categories in a conversion path, and in varying amounts, based on a first attribution model or algorithm identified as attribution model 0. The system may provide the attribution by calculating an amount of credit, for example denominated in generic units (e.g., a percentage of the 1 conversion event), in currency, or in other units. The system may store the attribution data associated with each category in a memory.

[0083] At blocks 212 and 214, the system may be configured to provide attribution of one or more of the same conversion events processed in block 210, though in this case using one or more additional different attribution models. Some exemplary attribution models will be described below, but may include such models as first click, first touch, category touch, recency, or combinations thereof. Additional sets of attribution data may be generated at

blocks 212 and 214 and stored in memory. In one implementation, blocks 212 and/or 214 may use attribution models other than a model based solely on the last click in the conversion path.

[0084] In one example, the conversion criteria themselves may be part of an attribution model, for example, if the attribution model includes a lookback window. Therefore, multiple conversion criteria may be applied as needed by different attribution models in the process. In one implementation, selection of conversion events (block 206) should be the broadest match selection, e.g., rules to define what counts as conversion for all attribution models to be used.

[0085] At block 216, the system may be configured to report the results of the different attribution models in textual, graphical, or other formats, some examples of which will be described below.

[0086] At block 218, the system may be configured to receive cost data representing a relative or actual cost of a plurality of different categories in a conversion path for a particular advertiser. The system may be configured to combine the attribution data sets with the cost data in any of a number of ways, for example by calculating ratios, percentages, etc. At block 220, the system is configured to generate report data based on a plurality of sets of attribution data and the cost data (e.g., return on ad spend or ROAS). The report data generated may be transmitted to a client device for display. The advertiser can then use the report data to make more informed decisions about the cost effectiveness of using certain advertising categories.

[0087] Types of attribution models

[0088] Referring now to FIG. 3, an illustration of conversion path data is shown, according to a described implementation. The illustration 300 (and similar illustrations in the other figures) may or may not represent graphical elements that would be part of report data transmitted to a client device. In this example, using the methods described above, the systems has stored a conversion path comprising six user interactions that led to a conversion event of, in this case, arriving at an advertiser site 302 and making a purchase at the website (either of which could be a conversion event in various embodiments). Each event or user interaction in the conversion path comprises category identifier data, which may identify the specific category and/or category type.

[0089] For example, category 304 is a display network channel. A display network channel may be a channel associated with a network of websites all signed up with or registered with a particular advertisement publishing network. The publishing network may be an entity

which buys specific placements on many websites and matches many different advertisements from different advertisers to the placements. In the case of category 304, an impression of an advertisement for advertiser site 302 or a product thereon was shown on May 15, 2011 at 7:32:21 PM.

[0090] A couple weeks later, on May 30, 2011 at 12:41:08 AM, the same user was provided with an impression for advertiser site 302 through a second category, different than the first category, namely a content-oriented website identified as Magazine. This impression is recorded when the user visits the website and stored for later use.

[0091] A while later, on June 11, 2011 at 3:25:02 PM, the user visits an affiliate blog category 308 and interacts with the category by clicking on the link associated with advertiser site 302. An affiliate category may be a price comparison-type website or other click tracker website which gets credit or compensation for a user clicking on information on the website. About 5 hours later, the user interacts with yet another different category, a webpage 310 showing an advertisement in an advertisement slot on the webpage.

[0092] The next day, the user performs a search using a search engine category 312. About six minutes later, a conversion event is finally detected at conversion event 314. The user has reached the advertiser website.

[0093] The advertisements viewed or otherwise interacted with in this example may be from one advertisement campaign or from multiple campaigns running at once, all associated with the advertiser and/or at least one same conversion goal.

[0094] A conversion event has happened, and the system is configured to attribute credit for the conversion event to one or more categories or events in the conversion path.

[0095] Referring to FIG. 4, an illustration of a last click attribution model is shown. Illustration 400 appears similar to illustration 300, except that a graph 402 has been overlaid on the conversion category illustrations to generally represent an amount of attribution credit to be allocated to each category. In the last click model, the system is configured to give all or substantially all of the credit (represented by bar 404) for the conversion to the last click in the conversion path before the conversion event, within a given time window. If no click is present in the time window, it assigns a credit to the closest display impression that is within the impression window.

[0096] Referring to FIG. 5, an illustration of a first click attribution model is shown. Illustration 500 shows the effect of running the first click attribution model on the conversion path data. In this case, the first click in the conversion path is attributed all or substantially all of the credit, as shown by bar 502. The first click model is effected in part by the

lookback window selected by the advertiser. Credit is given to the furthest active (click) type exposure to the conversion that is within a given window. If no click is present in the time window, it assigns a credit to the furthest away display impression that is within the impression window.

[0097] Referring to FIG. 6, an illustration of a first touch attribution model is shown. First touch rewards all or substantially all of the credit to the first event in the conversion path, regardless of whether it is a click, impression, or other event. In this case, the display network category is awarded all of the credit for the conversion, as indicated by bar 602. Credit is assigned to the furthest exposure away from the conversion that is within the lookback window. In another example, a first touch attribution model may be adjusted to determine the earliest point in time prior to conversion that a marketing category would be credited by this model, using a lookback window specific to this model.

[0098] Referring to FIG. 7, an illustration of a category touch attribution model is shown. The category touch model or reach model credits shares credit among multiple categories. In this case, all categories are given an equal portion of the credit for the conversion, in this case one-fifth of the credit to each of five categories. A variation on this model is to credit campaigns equally. For example, if campaign A represents two categories in the conversion path and campaign B represents one category in the conversion path, each of the campaign A and campaign B will receive an equal portion of the credit. A category may be defined for purposes of this model as a campaign or medium.

[0099] Referring to FIG. 8, an illustration of a recency attribution model is shown. The recency model rewards proximity to conversion on an ascending scale of credit. This model may further award more weighting of credit for clicks than for displays. Referring to illustration 800, the search click category 802 receives the largest portion of credit (as indicated by bar 804). The second largest portion of credit is allocated to the affiliate click category 806 (as indicated by bar 808). Although category 806 is further from the conversion than the display view category 810, category 806 is awarded a higher portion of the credit due to the click. Categories 812 and 814 are awarded the lowest amount of credit because they are furthest from the conversion and are not associated with any clicks.

[0100] The algorithms for calculating these attribution models may comprise parameters which can be adjusted, for example by the advertisers. For example, the recency weighting may be linear or exponential. The difference in weighting between clicks and views may have a multiplier which is adjustable.

[0101] In one particular example, in the recency model, credit may be assigned using a continuous exponential decay function that reflects the non-linear way that people remember or are influenced by exposure to information. This function may have two parameters: the decay half life of the function and an ‘active’ event type multiplier by which raw scores are multiplied. The function may also apply a logical test that discards conversions that have low total scores. An exemplary algorithm for such a model is provided below:

```
memoryDecay(report, conversionWindow, activeMultiplier, decayFactor)
{
    for (Conversion c : report)
    {
        conCur = {};
        /* scoring phase stores in conCur */
        for (Exposure e : c.exposures())
        {
            if (withinWindow (e, conversionWindow))
            {
                if (isClick(e))
                {
                    conCur.put(e=>activeMultiplier)
                }
                else
                {
                    conCur.put(e=>1.0);
                }
            }
            else
            {
                conCur.put(e=>
                    exp(-(interval(c, e) - conversionWindow) /
                        decayFactor));
                if (isClick(e))
                {
                    conCur.put(e=>conCur.get(e) *
                        activeMultiplier);
                }
            }
        }
        /* Logical test */
        total = 0;
        for (e=>score : conCur)
        {
            total = total + score;
        }
    }
}
```

```

    }
    score = 1 - (1/(1 + total));
    /*if low score then discard
    if (score<0.5)
    {
        for (e=>score : conCur)
        {
            conCur.put(e=>0.0);
        }
    }
    /* Record the scores for the exposure */
    for (e=>score : conCur)
    {
        recordScore(e, score);
    }
}
}

```

[0102] In one embodiment, the active multiplier may be calculated using a function which calculates the proportion of active conversions as a proportion of conversions:

$$100 / 1 - ((clickP + clickP) / (clickP + impressionP))$$

[0103] The value of clickP is the count of presence of conversions including clicks. The value of impressionP is the count of conversions including impressions. This function is predominantly driven by the presence of display-only conversion paths. The numerator includes two clickPs to balance the formula for the case that there are only click paths where all terms will be equal. The formula is biased toward clicks. For conversions which only include a click, a count of one may be added to impressions because there would be an impression if there was a click.

[0104] The equation shown above may provide an up weighting for clicks over impressions, where the impression weight is 100. For example, if this parameter is 200, then a click is valued at 2x an impression.

[0105] The decay factor used may be based on how fast customers convert. The decay factor may be calculated based on a time lag from first touch to conversion for all conversions. Then all of the conversions may be sorted by lag from first touch to conversion, to provide cumulative time to conversion. The time conversions may be listed and analyzed to reach the 50th percentile of the conversion base. This 50th percentile represents the time required for at least half of the fastest converters to achieve a conversion. A predetermined threshold of time (e.g., 2 hours) can be used so that when the time for the 50th percentile is less than 2 hours, 2 hours is used. This time is used as a “half-life” to conversion for the time

decay factor by calculating. The half life determines at what time interval the function allocates 0.5 raw un-weighted credits, as shown in the following equation:

$$(-halfLife / \log(0.5))$$

[0106] This gives a decay factor that results in 0.5 credit being assigned after 1 half life. Ultimately as this is driven by the conversions in the report it is essentially fit to the type of campaigns/media mix being used.

[0107] One justification for the recency attribution model is that as time passes, the influence of exposure to or interaction with marketing categories may decline due to memory effects, and the decay factor accounts for this. For this reason a recency model can be referred to as a 'memory' model and the curve as a 'memory' or 'forgetting' curve.

[0108] Time to conversion may be calculated in whole hours for a cumulative time to conversion. For example:

Time To Conversion Cumulative Percentage of Converters

0	20%
1	35%
2	45%
3	50%
4	55%

[0109] In this example, in less than one hour, 20% of users convert. Within 3 hours 50% of the users have converted. In this case, 3 hours is used as the half life.

[0110] For any of the models or functions described above, the system is configured to score or credit the categories as shown. Next, output data is grouped or accumulated. For example, all exposures or categories and their respective scores are examined and summed against a given group, determined by the category. For example, credit given to campaigns may be calculated by a function and then iterated over all recorded scores and then each exposure score added to the running total for the campaign to which the exposure belongs. This score accumulation process can be repeated for any dimension associated with the exposure including site name, advertiser, creative, placement, etc.

[0111] A final accumulation stage may involve calculating the percentage of credit for each grouping. This calculation may include summing the total credit for a given function, then dividing through all of the groups.

[0112] Another accumulation action may include factoring through revenue. The system may be configured to reweigh all function score totals per conversion to 1 and factor through

the conversion value. The result of this re-weighting is to change the distance of values for the domains of each conversion. The system may be configured to use a percentage view of results to yield the most consistent result. The system may then compare to spend percentage by the given dimension group.

[0113] Multiple attribution algorithms may be applied to the same user interactions (e.g., ad exposures, active clicks, conversion data, etc.) and the results compared.

[0114] Referring now to FIG. 9, an illustration of exemplary display data or report data showing results of multiple attribution model calculations will be described. In this example, credit has been attributed to multiple categories in a plurality of conversion paths for a marketing campaign for “Acme Insurance.” A first category 902 represents a search engine search for “Acme Insurance Insurance” keywords. A second category 904 represents an affiliate category. A third category 906 represents an online display campaign for Acme Insurance product A. A fourth category 908 represents a search engine search for “Car Insurance” keywords. A fifth category 910 represents an online display campaign for Acme Insurance product B.

[0115] The first (top) bar in each set of bars in the bar graph represents an amount of credit attributed to the category by the system using a recency attribution model. The second bar in each set of bars represents the amount of credit attributed using a last click attribution model. The third bar represents a campaign category distribution model. The fourth bar represents a first click attribution model, and the third bar represents a first touch attribution model. Each bar may be based on attribution data generated by the system.

[0116] The advertiser may make certain observations about the data. For example, First Click gives about the same credit as Last Click for each of the categories, as shown by arrows 912. Second, the recency model rewards affiliates slightly more than the last click model, as shown by arrows 914. Third, display has the greatest presence in the first touch and category models, as shown by arrows 916. The report data or display data may be grouped by site grouping, campaign grouping, creative level grouping, placement level grouping, or other groupings. The groupings should add up to 100% of the credit. As shown at 918, the grouping of this data is by campaign. Alternatively, the grouping may be site, sit placement, etc. Use of the terms category or marketing category herein may refer to categories on any of these levels.

[0117] Referring now to FIG. 10A, an illustration of exemplary display data or report data showing results of multiple attribution model calculations comprising cost data will be described. The system may be configured to generate cost data in any of a number of ways.

The cost data may be part of a return-on-investment or return on advertising spend analysis. The system may be configured to receive costs data representing a relative or actual cost of a plurality of categories in a conversion path. The analysis may be configured to compare the relative or actual cost to the credit given to each category. The cost data may comprise the media spend data for each category analyzed.

[0118] Referring to category 1 in this example, the received cost data may be presented in a bar graph form at bar 1002, indicated as 15% of the cost of the marketing campaign. The attribution data indicates at bar 1004 that, using attribution model B, category 1 receives 25% of the credit. The attribution data further indicates at bar 1006 that, using attribution model A, category 1 receives 20% of the credit. A conclusion may be drawn by an advertiser or the system that because credit percentage exceeds the media cost percentage using both attribution model algorithms, this category appears to be a relatively cost effective category.

[0119] Turning to category 2, media cost is 35%, model B allocates 30% of the credit and model A allocates 25% of the credit. A conclusion may be drawn by an advertiser or the system that because credit percentage is below the media cost percentage for both models, this appears to be a relatively expensive category.

[0120] Turning to category 3, media cost is 25%, model B allocates 20% of the credit for conversions to category 3 and model A allocates 30% of the credit to category 3. In this case, the credit percentage relationship to cost varies by model. The decision of whether the category is cost effective or expensive, and how to allocate future cost decisions depends on the advertiser or system's view on how the exposure path drives value.

[0121] As shown, the system may be configured to generate return on ad spend data based on credit or attribution data and cost data. The system may be configured to divide credit by the cost and present the data to a user in any of a variety of formats, such as textual, ratio, graphical, such as bar graph, or other formats. The system may further be configured to highlight or flag some data using different colors, highlighting, underlining, bolding, sounds, etc. For example, a credit to cost ratio of greater than 1 may be identified in the report data as efficient (e.g., with a green color), while a credit to cost ratio of less than 1 may be identified in the report data as inefficient (e.g., with a red color).

[0122] The system may be configured to perform a return on investment or return on advertising spend by receiving marketing spend and attributing credit of multiple models and calculating ratios of model credit to spend. The system may be configured to allocate the data over multiple sites, multiple campaigns, multiple creatives, etc.

[0123] Referring now to FIG. 10B, an illustration of exemplary display data or report data showing results of multiple attribution model calculations comprising cost data will be described. The data used in this illustration is the same as the data from FIG. 10B. In FIG. 10B, the cost data is presented in textual format comprising percentages and ratios, along with colored highlighting. In the case of category 1, a ratio of 1.33 “model A” credit/cost is presented in green text and a ratio of 1.67 “model B” credit/cost is presented in green text, because both models indicate an effective spend. In the case of category 2, both ratios are presented in red text. In the case of category 3, model A credit/cost ratio of 1.20 is presented in green text model B credit/cost ratio is presented in red text. Other colors or highlighting options are contemplated. The ratios may be presented in a matrix or table for ease of visualization.

[0124] FIG. 11 illustrates a depiction of a computer system 1100 that can be used to provide user interaction reports, process log files, implement an illustrative performance analysis apparatus 130, or implement an illustrative advertisement management system 110. Computing system 1100 includes a bus 1105 or other communication component for communicating information and a processor 1110 coupled to bus 1105 for processing information. Computing system 1100 also includes main memory 1115, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 1105 for storing information, and instructions to be executed by processor 1110. Main memory 1115 can also be used for storing position information, temporary variables, or other intermediate information during execution of instructions by processor 1110. Computing system 1100 may further include a read only memory (ROM) 1110 or other static storage device coupled to bus 1105 for storing static information and instructions for processor 1110. A storage device 1125, such as a solid state device, magnetic disk or optical disk, is coupled to bus 1105 for persistently storing information and instructions.

[0125] Computing system 1100 may be coupled via bus 1105 to a display 1135, such as a liquid crystal display, or active matrix display, for displaying information to a user. An input device 1130, such as a keyboard including alphanumeric and other keys, may be coupled to bus 1105 for communicating information, and command selections to processor 1110. In another embodiment, input device 1130 has a touch screen display 1135. Input device 1130 can include a cursor control, such as a mouse, a trackball, or cursor direction keys, for communicating direction information and command selections to processor 1110 and for controlling cursor movement on display 1135.

[0126] According to various embodiments, the processes that effectuate illustrative embodiments that are described herein can be implemented by computing system 1100 in response to processor 1110 executing an arrangement of instructions contained in main memory 1115. Such instructions can be read into main memory 1115 from another computer-readable medium, such as storage device 1125. Execution of the arrangement of instructions contained in main memory 1115 causes computing system 1100 to perform the illustrative processes described herein. One or more processors in a multi-processing arrangement may also be employed to execute the instructions contained in main memory 1115. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement illustrative embodiments. Thus, embodiments are not limited to any specific combination of hardware circuitry and software.

[0127] Alterations in calculations in attribution modeling based on feedback

[0128] Referring generally to FIGS. 12-15, the attribution models as described in the present disclosure may be used to provide feedback to adjust bids an automated bidding system. By adjusting the bids, a customer (e.g., an advertiser) may optimize their expenditures on advertisements. In the embodiments shown in FIGS. 12-15, the attribution models define an association between clicks, impressions, conversions, and other user activity related to the advertisements shown. These associations are then used to generate bids to be put on one or more keywords that are preferred by the attribution model.

[0129] Referring to FIG. 12, a flow diagram of a process 1200 for providing attribution based on advertisement conversion data in accordance with a described implementation. At a block 1202, process 1200 includes receiving user interaction data. The user interaction data may specify user interaction with content items and conversion items. A conversion item is a predetermined user action that satisfies a conversion criteria, as described above. At a block 1204, process 1200 includes determining that a conversion event has occurred based on the user interaction data and the conversion criteria. At a block 1206, process 1200 includes storing conversion path data based on user interaction data. The conversion path data includes user interaction data prior to and including the conversion event. Blocks 1202-1206 may generally resemble blocks 202-206 of the process of FIG. 2.

[0130] At a block 1208, process 1200 includes attributing a conversion event to a category in the conversion path using the attribution model to generate attribution data. The attribution model allocates credit for the conversion event to a plurality of different categories including a plurality of clicks in the conversion path (blocks 210-214 of Fig. 2). The

attribution model may be re-scaled and re-weighted using feedback before generating the attribution data. Such a process is shown in greater detail in FIG. 13.

[0131] At a block 1210, process 1200 includes generating a keyword bid based on the attribution model output. The keyword bid is a bid on a keyword placed by a customer. The bid placed or generated by the customer (e.g., an advertiser) represents a cost the customer is willing to pay for display of an advertisement associated with the keyword. An advertisement management system (e.g., advertisement management system 110 of FIG. 1) may then select the best keyword bid among all bids provided to the system. The selection of the keyword bid may be based on cost (e.g., the highest bid or highest amount of money for which an advertiser is willing to spend to have the ad displayed may be selected), expected performance of the ad (e.g., the system may estimate which advertisement may perform best and select the ad for display), or a combination of the two factors, or other factors.

[0132] An attribution model output may be used as feedback by the automated bidding system to adjust bids that are placed on keywords. The attribution model output is used to define the association between clicks, impressions, conversions, and other user activities related to advertisements. For example, in the example of Fig. 10A, the advertiser may bid on three keywords or sets of key words (keyword 1, keyword 2, and keyword 3) which are respectively related to the content in the three categories (category 1, category 2, category 3). In other words, when the advertiser wins bids relating to keyword 1, the advertiser's advertisement is shown in category 1; when the advertiser wins a bid relating to keyword 2, the advertiser's advertisement is shown in category 2; and when the advertiser wins a bid relating to keyword 3, the advertiser's advertisement is shown in category 3. (As will be appreciated, each category may also have multiple keywords associated therewith, e.g., a different keyword or set of keywords for each website in the category). In the example of Fig. 10A, category 1 has been determined to be a relatively cost effective category as compared to category 2. Accordingly, the automated bidding system may increase the bids for the keywords associated with category 1 and decrease the bids for the keywords associated with category 2 while staying within the overall budget limit set by the advertiser. Although the advertiser bids less for the keywords associated with category 2, the advertiser may still win at least some of the auctions relating to the keywords associated with category 2. Given that the new bids for category 1 are higher and the new bids for category 2 are lower, category 1 and category 2 may tend toward the same level of cost effectiveness in steady state, and a more optimal allocation of spending between category 1 and category 2 may be attained.

[0133] Process 1200 describes generating a keyword bid based on an attribution model output, and may be executed for any number of attribution models. Using these attribution model outputs, the bids to be submitted for one or more keywords may be changed. For example, referring also to FIG. 10A, according to a first attribution model, keywords in category 1 are associated with more conversions compared to the number of conversions associated with category 1 based on a second attribution model. This difference can be used as a signal to bid more or less on the keywords in category 1 based on the preference of which attribution model to use.

[0134] Further, a combination of several attribution models may be used instead of just using a difference between two attribution models as described above. For example, if several attribution models each indicate that the number of conversions associated with keywords in category 1 is lower than a desired threshold set by the advertiser, then the resulting outputs of the attribution models may be used as feedback, indicating that a bid on the keywords in category 1 should be reduced. If each attribution model indicates that the number of conversions exceed the threshold, then a bid on the keywords in category 1 should be increased. Similarly, any combination of outputs from the attribution models can be used to change the bids. It should be understood that the mathematical combination of outputs from the attribution models may vary without departing from the scope of the present disclosure.

[0135] Referring to FIG. 13, a flow diagram of a process 1200 for providing attribution based on advertisement conversion data in accordance with a described implementation. Fig. 13 is similar to Fig. 2 except that, in Fig. 13, other details are shown and it is assumed that user interaction data has previously been stored and that such data is to be used to update automated bidding parameters. Process 1300 may either be initiated by a customer request at block 1302 (e.g., when managing an advertising campaign) or an automated request at block 1304 (e.g., during real time of automated bidding system of advertisement management system 110 immediately prior to submission of a bid). The customer request (e.g., an advertiser request) of block 1302 may be initiated by a customer when the customer desires to optimize their expenditures on advertisements.

[0136] At a block 1306, process 1300 includes checking if conversions are tracked and receiving the conversion data if it is tracked. The conversion data may be historical conversion data relating to past conversions. The conversion data may include data provided by advertisers, according to one implementation. Such conversion data may include, for

example, sales information related to purchases made as a result of an advertisement. The conversion data may generally be the conversion data as described in FIGS. 1-11.

[0137] At a block 1308, attribution models are loaded (e.g., the attribution models as generally described in FIGS. 1-11). At a block 1310, model weights for the attribution models are received or generated. In one implementation, the model weights are received from a database or other source. For example, the weights may be provided by the advertiser. In another implementation, the model weights may be generated automatically. The generation of the model weights may be based on user preferences or other preferences. For example, there may be a preference to weigh a particular conversion event more than another conversion event (e.g., a “last click” being weighed more than a “first click”).

[0138] The model weights represent a value for a given action in a conversion path. For example, referring also to FIG. 4, if a last click attribution model is desired, the model weights may give a substantial weight, or all weight, to the last click in the conversion path. As another example, referring also to FIG. 5, if a first click attribution model is desired, the model weights may give a substantial weight, or all weight, to the first click in the conversion path. The model weights may reflect any combination of first clicks, last clicks, first touches, category touches, recency data, etc. In one implementation, model weights may be equal for each click in a conversion path. In another implementation, model weights may only give a weight for particular portions of a path. The model weights may be varied as described above.

[0139] At a block 1312, process 1300 includes running the attribution models with the conversion data loaded at block 1306. At a block 1314, the attribution models are then re-scaled. The models are re-scaled to ensure compatibility. The attribution models are re-scaled by scaling the metrics (e.g., the conversion data) used with the attribution models.

[0140] Re-scaling the models based on the conversion data enhances the accuracy of the model output. The original conversion data may not properly reflect a total amount of conversions since the conversion data may relate to different actions in a conversion path. For example, if the conversion data includes a total number of assisted conversions (e.g., conversions that were assisted by the presence of an impression or by a click from another website), such conversion data may be re-scaled to more properly reflect the total number of conversions that occurred. The re-scaled conversion data is then used in the attribution model to create the re-scaled attribution model. The process of re-scaling metrics (e.g., the conversion data) is shown in greater detail in FIGS. 14-15.

[0141] At a block 1316, the scaled conversion data from block 1314 is re-weighted. The re-weighting of the conversion data is based on the model weights received or generated at block 1310. After the re-scaling activity of block 1314, the model needs to be re-weighted based on the new conversion data. The re-weighting is shown in greater detail in FIGS. 14-15.

[0142] At a block 1318, process 1300 includes calculating cost data based on the re-weighted conversion data and the re-scaled attribution model. In one implementation, the calculation includes cost per action (CPA) calculations (e.g., a cost divided by attributed credit calculation). Such a calculation is used to determine how much a particular action is worth to an advertiser. Such a cost is used to determine how much the advertiser should bid on a keyword. In another implementation, the calculation includes return on investment (ROI) calculations (e.g., calculations where the modeled revenue of an advertisement (the amount of revenue an advertisement is expected to make for the advertiser) is divided by a cost of displaying the advertisement). Such a calculation is used to determine how much an advertiser is expected to receive in terms of revenue per action in a conversion path, and comparing it to a cost such as a keyword bid.

[0143] At a block 1320, process 1300 includes generating a keyword bid, report, or other output based on the cost data. In one implementation, a keyword bid is generated or calculated as described in process 1200 of FIG. 12. In another implementation, the cost data may be fed into another bid optimizing tool configured to generate a bid for an advertisement of the customer. In yet another implementation, a report may be generated based on the attribution data and cost data. For example, the report may include ROAS information, CPA information, return on investment (ROI) information, or other attribution data and cost data.

[0144] It may also be noted that the attribution models may also provide feedback for the calculations in block 1318. For example, the CPA metric is defined above as a total cost divided by total attributed conversions. Hence, the CPA metric may be impacted by attribution model outputs. This derived metric (CPA) can be used to adjust bids on keywords as well as the actual attribution model outputs. For example, the advertiser may have a target CPA, and bids may be adjusted to reach the target CPA.

[0145] Referring now to FIGS. 14-15, an example of re-scaling attribution models and re-weighting conversion data is shown in accordance with a described implementation. In FIG. 14, a table of conversion data is shown. The conversion data is historical conversion data that may be retrieved in a process for updating an attribution model (e.g., process 1300). In FIG. 14, conversion data is stored for two advertisements (a generic ad and a branded ad).

Conversion data relating to the number of impression-assisted conversions the ad triggered, the number of first clicks of the ads, the number of click-assisted conversions the ad triggered, and the number of last clicks of the ads are included in the historical data. Also included in the table is a cost associated with each ad, which represents the cost the advertiser is willing to spend to have the ad chosen and displayed by an advertisement management system.

[0146] Referring now to FIG. 15, the historical conversion data is re-scaled. For example, the number of total assisted conversions (200) for the generic advertisement is selected and compared to the total number of conversions (400). The total number of conversions, in this example, is represented by the number of last clicks on the advertisement. Therefore, an adjustment factor (e.g., an amount to scale by) is defined as the number of total conversions divided by the total assisted conversions. In this example, the adjustment factor is 2. This adjustment factor is then used to adjust all subsets of the assisted conversions.

[0147] Referring further to FIG. 15, all of the total assisted conversions were click-assisted conversions. Therefore the number of click-assisted conversions (200) is multiplied by the adjustment factor (2) to obtain an adjusted click-assisted conversion total of 400. As another example, if there were 100 click-assisted conversions, the adjusted click-assisted conversion total would be 200. This adjusted click-assisted conversion is then used in the attribution model instead of the original number of click-assisted conversions.

[0148] As an example, say that there are three metrics (e.g., three types of conversion data) that are to be used in an attribution model. The re-weighting of the attribution model may consist of the following equations:

$$\begin{aligned} &\text{Metric A} * \text{Factor A} * \text{Model Weight A} \\ &\text{Metric B} * \text{Factor B} * \text{Model Weight B} \\ &\text{Metric C} * \text{Factor C} * \text{Model Weight C} \\ &= \textit{Attributed Conversions} . \end{aligned}$$

[0149] In this example, the metrics A, B, C are the conversion data, the factors A, B, C are the adjustment factors described above, and the model weights A, B, C, are weights loaded for the attribution model that control how much impact each metric has on the attribution model. The model weights of the attribution model should add up to 100%. After the attributed conversions are calculated using the above equations, cost data may be calculated and keyword bids may be generated as described above.

[0150] **Configurations of various exemplary implementations**

[0151] Implementations of the subject matter and the operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Implementations of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions, encoded on one or more computer storage medium for execution by, or to control the operation of, data processing apparatus, such as a processing circuit. A processing circuit may include any digital and/or analog circuit components configured to perform the functions described herein, such as a microprocessor, microcontroller, application-specific integrated circuit, programmable logic, etc. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. A computer storage medium can be, or be included in, a computer-readable storage device, a computer-readable storage substrate, a random or serial access memory array or device, or a combination of one or more of them. Moreover, while a computer storage medium is not a propagated signal, a computer storage medium can be a source or destination of computer program instructions encoded in an artificially-generated propagated signal. The computer storage medium can also be, or be included in, one or more separate components or media (e.g., multiple CDs, disks, or other storage devices). Accordingly, the computer storage medium is both tangible and non-transitory.

[0152] The operations described in this specification can be implemented as operations performed by a data processing apparatus on data stored on one or more computer-readable storage devices or received from other sources.

[0153] The term “data processing apparatus” or “computing device” encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, a system on a chip, or multiple ones, or combinations, of the foregoing. The apparatus can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can also include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, a cross-platform runtime environment, a virtual machine, or a combination of one or more of them.

The apparatus and execution environment can realize various different computing model infrastructures, such as web services, distributed computing and grid computing infrastructures.

[0154] A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, object, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub-programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0155] The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform actions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA or an ASIC.

[0156] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a processor for performing actions in accordance with instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device (e.g., a universal serial bus (USB) flash drive), to name just a few. Devices suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example

semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0157] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending webpages to a web browser on a user's client device in response to requests received from the web browser.

[0158] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), an inter-network (e.g., the Internet), and peer-to-peer networks (e.g., ad hoc peer-to-peer networks).

[0159] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some embodiments, a server transmits data (e.g., an HTML page) to a client device (e.g., for purposes of displaying data to and receiving user input from a user interacting with the client device). Data generated at the client device (e.g., a result of the user interaction) can be received from the client device at the server.

[0160] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any inventions or of what may be claimed, but rather as descriptions of features specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0161] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0162] Thus, particular implementations of the subject matter have been described. Other implementations are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results. In addition, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

WHAT IS CLAIMED IS:

1. A computer system for providing attribution based on advertisement conversion data, the system comprising:
 - a processing circuit configured to:
 - receive user interaction data, wherein the user interaction data specifies user interaction with content items and conversion items, wherein a conversion item is a predetermined user action that satisfies a conversion criteria;
 - determine that a conversion event has occurred based on the user interaction data and the conversion criteria;
 - store conversion path data based on the user interaction data, wherein the conversion path data comprises user interaction data prior to and including the conversion event;
 - attribute the conversion event at least in part to a category in the conversion path using an attribution model to generate attribution data, wherein the attribution model allocates credit for the conversion event to a plurality of different categories including a plurality of clicks in the conversion path; and
 - generate a marketing purchase parameter based on output of the attribution model.
2. The computer system of Claim 1, wherein the attribution model comprises weights determining how credit is allocated for the conversion event between the plurality of different categories.
3. The computer system of Claim 2, wherein the processing circuit is further configured to receive weights for the attribution model via a user interface.
4. The computer system of Claim 2, wherein the processing circuit is further configured to generate weights for the attribution model.

5. The computer system of Claim 1, wherein the attribution model allocates more credit to an event in the conversion path closer to the conversion event than a comparable event in the conversion path further from the conversion event.
6. The computer system of Claim 2, wherein the attribution model is updated based on feedback; wherein the updating comprises running the attribution models against conversion event history and re-scaling the attribution models to ensure compatibility with the conversion event history.
7. The computer system of Claim 6, wherein the attribution model updating is initiated based on a customer request; wherein the customer is an advertiser.
8. The computer system of Claim 6, wherein the attribution model updating is initiated based on an automated request from a bidding tool; wherein the bidding tool generates a marketing purchase parameter for a customer.
9. The computer system of Claim 6, wherein the attribution model is re-weighted after the attribution model is re-scaled.
10. The computer system of Claim 1, wherein the marketing purchase parameter is one of a keyword bid, a target cost, and a maximum cost.
11. The computer system of Claim 10, wherein the marketing purchase parameter is a target cost and wherein the target cost is one of a target cost per action, a target cost per mille, and a target cost per click.
12. The computer system of Claim 1, wherein the marketing purchase parameter relates to at least one of placement and display of an advertisement.
13. A method for providing attribution based on advertisement conversion data, comprising:

receiving user interaction data, wherein the user interaction data specifies user interaction with content items and conversion items, wherein a conversion item is a predetermined user action that satisfies a conversion criteria;

determining that a conversion event has occurred based on the user interaction data and the conversion criteria;

storing conversion path data based on the user interaction data, wherein the conversion path data comprises user interaction data prior to and including the conversion event;

attributing the conversion event at least in part to a category in the conversion path using an attribution model to generate attribution data, wherein the attribution model allocates credit for the conversion event to a plurality of different categories including a plurality of clicks in the conversion path; and

generating a marketing purchase parameter based on output of the attribution model.

14. The methods of Claim 13, wherein the attribution model comprises weights determining how credit is allocated for the conversion event between the plurality of different categories.

15. The method of Claim 14, further comprising receiving weights for the attribution model via a user interface.

16. The method of Claim 14, further comprising generating weights for the attribution model.

17. The method of Claim 13, wherein the attribution model allocates more credit to an event in the conversion path closer to the conversion event than a comparable event in the conversion path further from the conversion event.

18. The method of Claim 13, further including updating the attribution model based on feedback; wherein the updating comprises running the attribution models against conversion

event history and re-scaling the attribution models to ensure compatibility with the conversion event history.

19. The method of Claim 18, wherein updating the attribution model is initiated based on a customer request; wherein the customer is an advertiser.

20. The method of Claim 18, wherein updating the attribution model is initiated based on an automated request from a bidding tool; wherein the bidding tool generates the marketing purchase parameter for a customer.

21. The method of Claim 18, further including re-weighting the attribution model after the attribution model is re-scaled.

22. One or more computer readable media having instructions stored therein, the instructions being executable by one or more processors to cause the one or more processors to perform operations, comprising:

receiving user interaction data, wherein the user interaction data specifies user interaction with content items and conversion items, wherein a conversion item is a predetermined user action that satisfies a conversion criteria;


determining that a conversion event has occurred based on the user interaction data and the conversion criteria;

storing conversion path data based on the user interaction data, wherein the conversion path data comprises user interaction data prior to and including the conversion event;

attributing the conversion event at least in part to a category in the conversion path using an attribution model to generate attribution data, wherein the attribution model allocates credit for the conversion event to a plurality of different categories including a plurality of clicks in the conversion path; and

generating a marketing purchase parameter based on output of the attribution model.

23. The computer readable media of Claim 22, wherein the attribution model comprises weights determining how credit is allocated for the conversion event between the plurality of different categories.






FIG. 2




FIG. 3




FIG. 4




FIG. 5

600

602							
Display network	05/15/11	07:32:21 PM					Display network
Magazine	05/30/11	12:41:08 AM					Display view
Affiliate blog 1	06/11/11	03:25:02 PM					Affiliate click
Webpage	06/11/11	08:21:09 PM					Display view
Google	06/12/11	10:25:05 PM					Search click
CONVERSION	06/12/11	10:32:57 AM					Advertiser site

FIG. 6




FIG. 7

800




FIG. 8




FIG. 9






FIG. 10A

Media category	Model A credit %	Model B credit %	Media cost %	ROI ratio A credit/cost	ROI ratio B credit/cost	Decision
Category 1	20%	25%	15%	1.33	1.67	Consider higher spend
Category 2	25%	30%	35%	0.71	0.86	Consider reduced spend
Category 3	30%	20%	25%	1.20	0.80	Spend decision depends on taking a view on how the exposure path drives value

FIG. 10B






FIG. 13

Campaign	Cost	Impression-Assisted		Click-Assisted		Last Clicks (= Total)
Source	Ad Words	Search Funnels	First Clicks	Search Funnels	Search Funnels	Search Funnels
e.g. Generic	£500	462	24	200	400	
Branded campaign	£700	851	12	60	120	

FIG. 14

Total			Click-Assisted		
Total Assisted Conversions	Conversions	Adjustment factor	Campaign	Conversions	Adjusted CAC
200	400	$400/200 = \underline{2}$	e.g. Generic	200	$200 \times \underline{2} = 400$
			...		

FIG. 15

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/023186**A. CLASSIFICATION OF SUBJECT MATTER****G06Q 30/02(2012.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06Q 30/02; G06Q 30/00; G06F 15/18; G06Q 30/08

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & keywords: conversion, attribution model, advertisement

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2011-0071900 A1 (ANIL KAMATH et al.) 24 March 2011 See abstract, claims 1-19 and figures 1-9.	1-23
Y	US 2008-0275757 A1 (ABHINAY SHARMA et al.) 06 November 2008 See abstract, claims 1-10 and figures 2-4.	1-23
Y A	US 2011-0040636 A1 (WILLARD L. SIMMONS et al.) 17 February 2011 See abstract, claims 1-12 and figures 1A-7.	6-9, 18-21 1-5, 10-17, 22-23
A	KR 10-2010-0042965 A (NHN BUSINESS PLATFORM CO., LTD.) 27 April 2010 See abstract, claims 1-6 and figures 2-3.	1-23
A	KR 10-2009-0063448 A (NHN BUSINESS PLATFORM CO., LTD.) 18 June 2009 See abstract, claims 1-7 and figures 1-5.	1-23

☐ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

02 May 2013 (02.05.2013)

Date of mailing of the international search report

03 May 2013 (03.05.2013)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan
City, 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

OH, Eung Gie

Telephone No. 82-42-481-8744



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/023186

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0071900 A1	24.03.2011	US 2004-0167816 A1 US 2011-0087542 A1 US 7870017 B2 WO 2011-035156 A2 WO 2011-035156 A3	26.08.2004 14.04.2011 11.01.2011 24.03.2011 21.07.2011
US 2008-0275757 A1	06.11.2008	AU 2008-248091 A1 CA 2686407 A1 CN 101689273 A EP 2156389 A2 EP 2156389 A4 JP 2010-529523 A WO 2008-137194 A2 WO 2008-137194 A3	13.11.2008 13.11.2008 31.03.2010 24.02.2010 02.02.2011 26.08.2010 13.11.2008 30.12.2009
US 2011-0040636 A1	17.02.2011	CN 102576436 A EP 2465086 A2 US 2011-0040611 A1 US 2011-0040612 A1 US 2011-0040613 A1 US 2011-0040635 A1 WO 2011-020076 A2 WO 2011-020076 A3	11.07.2012 20.06.2012 17.02.2011 17.02.2011 17.02.2011 17.02.2011 17.02.2011 28.04.2011
KR 10-2010-0042965 A	27.04.2010	None	
KR 10-2009-0063448 A	18.06.2009	WO 2009-078591 A1	25.06.2009