Interactive Campaign Planning for Marketing Analysts

Fan Du
University of Maryland
College Park, MD, USA
fan@cs.umd.edu

Sana Malik
Adobe Research
San Jose, CA, USA
sana.malik@adobe.com

Eunyee Koh
Adobe Research
San Jose, CA, USA
eunyee@adobe.com

Georgios Theocharous
Adobe Research
San Jose, CA, USA
theochar@adobe.com

CHI'18 Extended Abstracts, April 21–26, 2018, Montreal, QC, Canada.
ACM 978-1-4503-5621-3/18/04.
https://doi.org/10.1145/3170427.3188531

Abstract
Event sequence data is generated across nearly every domain from social network activities and online clickstreams to electronic health records and student academic activities. Patterns in event sequences can provide valuable insights to assist people in making important decisions, such as business strategies, medical treatments, and careers plans. EventAction is a prescriptive analytics tool designed to present and explain recommendations of event sequences. EventAction provides a visual and interactive approach to identify similar records, explore potential outcomes, review recommended action plan that might help achieve the users' goals, and interactively assist users as they define a personalized action plan. This paper presents the first application of EventAction in the digital marketing domain. Our direct contributions are: (1) a report on two case studies that evaluate the effectiveness of EventAction in helping marketers prescribing personalized marketing interventions and (2) a discussion on four major challenges and our solutions in analyzing customer records and planning marketing interventions.

Author Keywords
Digital marketing; web history; recommender systems.

ACM Classification Keywords
H.5.2 [User Interfaces]: Graphical user interfaces (GUI)

Introduction

Event sequence data has become ubiquitous with the development of mobile devices, electronic communication, and sensor networks. It can be collected everywhere from social network activities and online clickstreams, to electronic health records and student academic activities. Sequence recommender systems have been developed to assist people in making decisions by analyzing patterns in event sequence data, for example, recommending a sequence of places to visit in a park based on the trajectories of past visitors or recommending a series of marketing interventions to promote sales based on previous successful intervention strategies.

Most existing sequence recommender systems act like black boxes, not providing insight into the system logic or offering justification for the recommendations. Using these black-box systems, users usually have very limited knowledge about how the recommendations are generated, which may impair their confidence in following the recommended plan and discourage their engagement in the decision process. While such black-box techniques have been effective and successful in entertainment scenarios, previous research found that users are willing to spend more effort and want to be more engagement when making important decisions [2].

The main novelty of the approach of EventAction is that it properly presents and explains the recommendations to users, which is critical to the effectiveness of recommender systems and decision support tools in general. EventAction provides a visual analytics approach to (1) find similar archived records, (2) explore potential outcomes, (3) review recommended action plan that might help achieve the users' goals, and (4) interactively assist users as they develop a personalized action plan associated with a probability of success. In this paper, we present the first application of EventAction in the digital marketing domain. Our direct contributions are:

- A report on two case studies that evaluate the effectiveness of EventAction in helping marketers prescribing personalized marketing interventions.
- A discussion on the major challenges and our solutions in analyzing customer records and planning marketing interventions.

Overview of EventAction

This section provides a brief overview of EventAction's interface components (Figure 1). Complete descriptions and algorithmic details can be found in previous work [1, 2].

Record timeline: When using EventAction, analysts start by retrieving a seed record from the database. The activities of the seed record are shown in a time table, where each row is an event category and each column is a period of time (Figure 1a). Activities within each time period (e.g., one week) are aggregated and represented by a gray square. The sizes of the squares encode the numbers of event occurrences. Users can shorten the length of the time period for fine-grain analyses. Individual timelines of similar records are also displayed in time tables (Figure 1e). Timelines can be aligned by the first events or a specific time.

Similarity criteria controls: Each criterion is displayed as a rectangular glyph (Figure 1b) showing its name, the value of the seed record, and the value distribution of all archived records. The usage of a criterion can be switched among "Ignore", "Close Match", or "Exact Match". Users can define a tolerance range and treat values within the range as equivalent of the value of the seed record. Each criterion is associated with an adjustable weight. Users can add new

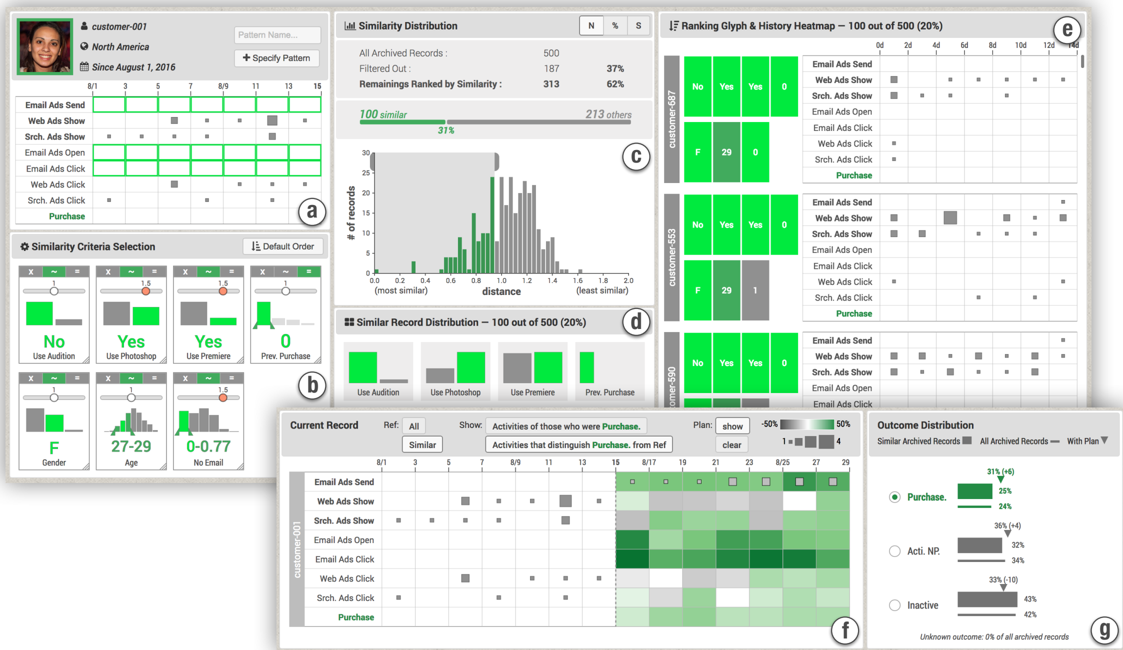

Figure 1: EventAction provides a data-driven approach for developing action plans to achieve the desired outcome. Continued on the left.

Figure 1 (cont'd). EventAction consists of seven coordinated views: (a) seed record timeline, (b) similarity criteria control panel, (c) similarity score distributions, (d) similar record distributions, (e) similar record timelines, (f) activity summary view, and (g) outcome distribution view.

This figure illustrates a synthetic dataset of a seed customer and 500 archived customer records. Marketing activities are related to email ads, web ads, and search ads. Record attributes include the customers' genders, ages, and previous product usages. Three types of outcomes are defined: "Purchase", "Active but No Purchase", and "Inactive".

All record attributes are used as similarity criteria by default (b) and a new criterion is created to capture the temporal pattern of having no email-related activities (a). The top 100 most similar records are selected as the peer group (c). An action plan of sending the customer more email ads is specified (f) and the likelihood of purchase increases by 6% (g).

criteria by specifying temporal patterns on the seed record timeline (Figure 1a).

Similar record distributions: After defining the similarity criteria, EventAction computes a similarity score for each archived record. A histogram of the scores (Figure 1c) is shown to help users specify a threshold and select a subset of the archived records as the peer group. The criteria values of the similar records will be summarized in barcharts to show the distributions (Figure 1d).

Outcome distribution: The outcome distribution view (Figure 1g) shows the outcome distributions of the similar records (thicker bars) and all archived records (thinner bars). From this view, users can estimate the most likely outcome of the seed record, the likelihood of achieving the desired outcome, and whether the seed record's likelihood of achieving the desired outcome is above or below the baseline of all archived records. Users can change the desired outcome during the analysis.

Activity summary and action plan: A summary of the activities of the similar record is integrated in the timeline of the seed record (Figure 1f). The darker the background color, the more popular this activity is in this time period. Users can select to show only the activities of records having the desired outcome and explore patterns that can be used to guide the specification of the action plan.

On top of the activity summary view, users can specify an action plan by adding events to the time table. As the plan is being made, EventAction will update outcome estimation incorporating the planned events into the seed record's timeline, giving users immediately feedback on how the plan affects the outcome likelihoods. EventAction can also automatically recommend a plan based on the activities of the similar records.

Case Studies

We report on two case studies conducted with 5 marketing analysts and using real-world event sequence datasets. Two of the analysts focused on email campaigns, two on cross-channel marketing, and one on web analytics. Each case study lasted about a month consisting of interviews, data preparation, system deployment, and data exploration. We investigated how EventAction can help marketers prescribing personalized marketing interventions.

Study 1: Customer Onboarding

In this case study, the analysts wanted to make plans for sending onboarding emails to new customers so as to increase their engagement.

Dataset: The analysts provided a dataset of 25,000 archived records of past customers who have received a series of 5 onboarding emails. The content of the emails including welcome notes, resource links, tutorials, and trial promotions. The dataset contains about 112,000 events tracking the send, open, and click (e.g., link clicks) of each email. We used a sample of 500 records and 8,000 events in the case study. Only one record attribute existed in the dataset indicating the regions of the customers. The outcome was defined by the number of emails a customer clicked ("0 click", "1-2 clicks", and "3-5 clicks").

Process: The analysts selected a seed record who have received and opened the first two emails but have not clicked any links. They wanted to make a plan for the subsequent emails that may lead to the outcome of "3-5 clicks." They started by specifying a "no click" pattern and only keeping customers having this pattern. Then, they selected the top 30% most similar records as the peer group and continued to review guidance for planning.

The analysts opened the activity summary view to review the email sending patterns of all archived records. The heatmap showed hotspots approximately every 7 days with some variations, which was expected by the analysts. From the outcome distribution view, the analysts realized that the seed record's likelihood of clicking 3-5 emails was only about 3%, which was much worse than the baseline of all archived records. The analysts decided to lower their expectations and changed the desired outcome to "1-2 clicks."

Then, they reviewed activities that distinguish customers who had "1-2 clicks" from others in the peer group. A green hotspot for email #3 showed up three days after sending email #2. About 11% more similar customers who received email #3 on that day will make 1-2 clicks during the onboarding. If they also open that email, the difference will further increase to 14%. The analysts checked the content of email #3 and found that it was featuring leaning resources and tutorials for the product. They explained: *"we thought it might be an important email and now EventAction provides evidence for it."* Following these findings, the analysts specified a plan for sending the subsequent emails. EventAction estimated an 11% increase in the seed record's likelihood of making 1-2 clicks.

Study 2: Channel Attribution Analysis

In this case study, the marketing analysts wanted to understand which campaign channels will be the most effective for converting a current customer into sales qualified.

Dataset: The analysts prepared a dataset of 997 archived records of past customers. The record attributes included which product was promoted and the region of the campaign. Campaign activities included “event invitation”, “paid search ads”, and “email sent”. Customers’ activities included “email open”, “email click”, and “website visit”. The outcome was defined by whether or not a customer became sales qualified judged by the sales team.

Process: The analysts selected a seed record who actively opened emails but never visited any product websites during the past 5 months. They reviewed the profile of the customer and found that their past interactions with this customer were mainly by email with only a few “event invitations” and no “paid search ads.” They created a new similarity criterion to reflect this pattern and selected the top 20% most similar records as the peer group.

The analysts immediately noticed that in the following 5 months those similar customers usually continue to actively receive and open emails. Their likelihood of becoming sales qualified was slightly below the baseline but still promising. The analysts switched to show activities distinguishing those who became sales qualified from others. Green hotspots showed up in the 6th and 7th months for “event invitation”, “email sent”, and “email click” indicating that sending out event invitations and campaign emails soon may help improving the outcome. The analysts specified a plan using these insights and the estimated likelihood increased by 10% which outperformed the baseline.

Feedback

Pseudo A/B testing: In both case study, the marketing analysts found EventAction useful for testing hypotheses based on historical data. They commented that EventAction allowed them to simulate plans and get results immediately, which can help selecting variables for A/B testings.

Temporal information: All marketing analysts liked EventAction’s visual and interactive way for exploring the temporal information as one said *“I can see the data directly.”* The analysts in *Study 2* also applaud that EventAction introduced a new time dimension for their attribution analysis because it not only informed them about which channels were important but also showed how the importance evolves over time. In addition, EventAction enabled them to filter the records using temporal patterns, which helps getting more precise results.

Automatic planning: The analysts were excited about EventAction’s automatic plan recommendation feature because *“it will save a lot of time and effort in the long term.”* However, they prefer to learn more about the mechanism before relying on it in real tasks. They suggested a workflow of showing the recommended plan at the beginning and allowing users to modify it during the analysis.

Challenges and Solutions

Through the process of the two case studies, the analysts have highlighted the challenges in analyzing customer records and planning marketing interventions. These challenges lie in both the uniquenesses of customer records and specific marketing tasks. We cover the 4 major challenges and discuss our solutions.

Challenge 1: Limited Record Attributes

Unlike patient or student records, customer records are usually anonymous without details such as demographics, diagnoses, or surveys. The available record attributes are usually very limited which makes it difficult to profile the customers and design personalized campaign strategies. EventAction addresses this challenge by using customer's activity patterns to identify similar customers and guide the planning. For example, given a customer who opens campaign emails but never visits the product website, marketers can find similar customers having this activity pattern and explore what campaign strategies worked the best for them.

Challenge 2: Visualizing Complex Temporal Data
Temporal data in the marketing domain are difficult to visualize due to their complexities in three aspects: (C2.1) the number of event categories is large capturing various campaign-related activities; (C2.2) the amounts of events in categories are very different, ranging from hundreds of email sents to only one or two purchases; (C2.3) many events occur at the roughly same time causing severe overlaps and visual clutters.

EventAction's timeline view can effectively handle event co-occurrences (C2.3) by aggregating events in each time period. However, since it uses the sizes of the squares to show the numbers of events, popular categories will dominate the view (C2.2), making squares in minor categories invisible. We addressed this issue by using a power scale $size = sqrt(num)$ when the range of the sizes is large. We also grouped the event categories into three classes to help users focus on one group at a time (C2.1): interventions, reactions, and outcome. However, a more scalable timeline design is still needed to fully address C2.1.

Challenge 3: Large Number of Records
A marketing dataset may contain millions of customer records, which can significantly slow down the computation and rendering. EventAction mitigated this issue by only visualizing similar records. To accelerate the similarity computation, we will investigate other techniques such as clustering and comparing records in groups. We will also replace the "seed record" with a seed group so that marketer can create a plan for a whole group of customers sharing common attributes or temporal patterns.

Challenge 4: Slow and Expensive A/B Testing
Conducting A/B testings to examine different campaign strategies may cost significant resources and take a long time when the number of variables is large. EventAction provides a low-cost approach allowing marketers quickly simulate different plans using historical data and get immediate results. The actual A/B testing will only need to cover strategies with promising results or low confidences (e.g., very few archived records matched the criteria).

Conclusion and Future Work
This paper presented the first application of EventAction in the marketing domain. We reported on two case studies evaluating the effectiveness of EventAction and discussed four major challenges and our solutions. In future work we will design scalable techniques to support analyzing and visualizing larger numbers of records and event categories. We will also investigate how to use automatic planning to facilitate users and inspire users' confidence.

REFERENCES
1. Fan Du, Catherine Plaisant, Neil Spring, and Ben Shneiderman. 2016. EventAction: Visual Analytics for Temporal Event Sequence Recommendation. In *Proceedings of the IEEE VAST*. 61–70.

2. Fan Du, Catherine Plaisant, Neil Spring, and Ben Shneiderman. 2017. Finding Similar People to Guide Life Choices: Challenge, Design, and Evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 5498–5509.