February 21, 2018

Conversion funnel optimization using machine learning

Gang Wang

Yiran Mao

Conversion funnel optimization using machine learning

ABSTRACT

Conversion funnels represent the journey of a consumer through a marketing campaign resulting in a sale. The described techniques incorporate marketer inputs and apply machine learning to optimize conversion funnels in online advertising. The techniques incorporate initial inputs from marketers and apply machine learning techniques to optimize conversion funnels in the context of advertising. The initial inputs include definitions of states and optimal actions to move consumers to next states of a funnel. With user permission, a machine learning model, e.g., that uses reinforced learning is randomly applied to try other actions defined by marketers (e.g., for other states) to collect training data on various actions. The training data is used to train the model and eliminates the requirement of a large amount of training data to be provided at the beginning of the process. The trained model thus obtained predicts optimal actions to move the consumer through the conversion funnel. Predicted actions using the model for a given state, along with supporting evidence, are provided to marketers for review and approval. Marketers can modify the funnel parameters, e.g., states and actions, based on such evidence.

KEYWORDS

- Conversion funnel
- Online advertising
- Reinforced learning
- Cold-start problem
- Action prediction
- Training data

BACKGROUND

A conversion funnel typically includes a series of states (also known as stages) and measurable criteria to evaluate whether a consumer is in a certain state. Advertising methods and messages are provided to consumers at each state to move consumers to the next state. Various methods (e.g., survey, checkout) are used to evaluate effectiveness of a conversion funnel and to optimize the funnel, for example, based on past experience, marketing theories, etc.

The design, execution, and optimization of a conversion funnel is largely based on empirical knowledge and past experience, and less on systematic data analysis and insights. Marketing researchers have developed extensive theories and best practices for designing, executing, and optimizing conversion funnels, also known as purchase funnels.

Existing online advertising products and services may not incorporate the concept of conversion funnel into ad campaigns. Even when campaigns indirectly apply the conversion funnel techniques, end to end optimization of the conversion funnel is not performed. Instead, online advertisers usually define disjoint ad campaigns, one for each state of the funnel, and thus, online ad products tend to optimize ad campaigns for different stages of the funnel separately.

DESCRIPTION

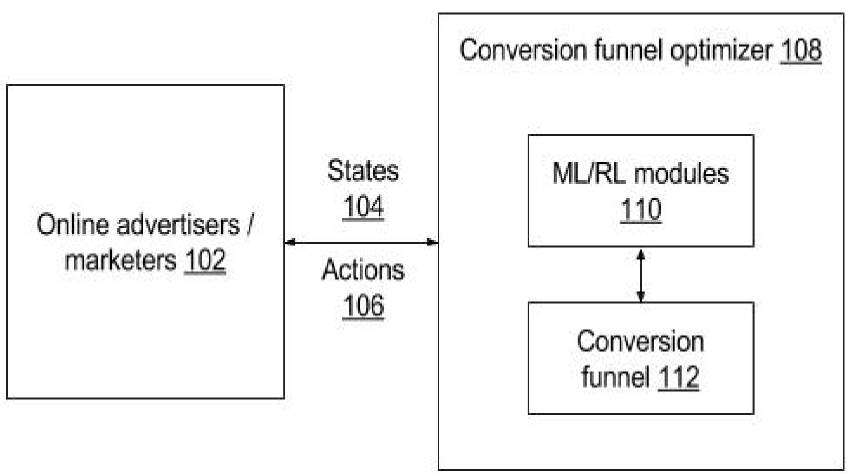

Fig. 1: Conversion funnel optimization

Referring to Fig. 1, techniques described enable online advertisers and marketers (102) to perform optimization of a conversion funnel by application of machine learning (ML). A computing device implements a conversion funnel optimizer (108) that includes Machine Learning or Reinforced learning (ML/RL) modules (110) to optimize conversion funnels (112). Marketers initially define the states (104) of a funnel, e.g., using best practices and organizational knowledge, and actions (106) at each state, along with the cost of each action, to optimize business objectives. For instance, sending a $200 ad coupon with an expected redemption rate of 80% to a consumer is an action that costs the marketer $160 for converting the sale, i.e., the objective.

By incorporating various Machine Learning/Reinforced learning (ML/RL) algorithms, machine learning (ML) modules can automatically predict actions that are most likely to lead to transitions between the different funnel states, and optimize the long-term reward. One way[1] to define the long-term reward can be maximizing the number of consumers moving through the funnel to the end state, e.g., making a purchase or completing a transaction. To optimize the conversion funnel, valid actions of the ML/RL models can be defined, e.g., showing consumers specific ads, sending promotional emails, asking consumers to create an account, all executed with express user permission and consent.

ML/RL algorithms within the ML modules are trained to select an action for each state with the goal of optimizing specific marketing objectives. For instance, the objective of increasing conversions can be achieved by moving more consumers to the end state within a

[1] There are other ways to define the long-term reward, for example, total value created minus total advertising cost.

given timeframe, while time taken by consumers to reach the end state can be reduced to reach the objective of optimizing the velocity of conversions.

Some of the marketing objectives are defined as constraints. Reducing total cost to generate a specified number of conversions within a fixed timeframe is an example of a constraint. Increasing conversions within a given timeframe and specified budget but with average cost per conversion not exceeding a predetermined amount is another constraint.

As the underlying machine learning model does not have any training data initially, for consumers at a given state, with a high probability P (e.g., P=0.9), the techniques perform an action specified by marketers, initially without using machine learning optimization. With a probability of (1-P), during a training phase, the model randomly tries other actions (e.g., specified in another state) to collect training data points in order to train the model and improve the accuracy of predictions from the model.

For a specific consumer at state S_m, if action A_n is taken, the consumer moves to state S_x with cost $C_{m,n,x}$ and probability $P_{m,n,x}$. By representing the states and consumer actions as a graph, for a consumer at state S_{start}, the most suitable action to take in order to move the consumer to the end state S_{end} with the minimum cost is predicted by the model. This is the classic "shortest path between two nodes in a graph" problem and can be solved using a deterministic and efficient algorithm when the ML can predict $C_{m,n,x}$ and $P_{m,n,x}$ accurately after sufficient training.

As the model trains on more data, the quality of the machine learning model improves, thereby increasing the confidence of the predicted action for each state. For example, data is collected to predict the state to which a consumer belongs and the cost of actions in each state, e.g., the cost of showing an ad. With permission from consumers for use of additional factors,

factors such as consumer demographic (e.g., gender and age group) data and information about consumer interests (e.g., history of online searches) can also be analyzed to better target marketing campaigns. The predicted optimal action for each state is provided to marketers for review along with supporting evidence. Marketers can then approve the predicted action and/or modify conversion funnel parameters, i.e., states or actions.

Optionally, or in association with marketer review, the techniques can also automatically modify the probability of below actions based on estimated quality of the corresponding graph.

- The action defined by marketers initially or the predicted action approved by marketers.
- The optimal action determined based on the graph but without marketer approval.
- Other marketer or predicted actions randomly chosen at a given state of the funnel.

The disclosure also addresses existing challenges for applying machine learning (ML) and reinforced learning (RL) techniques to optimize the conversion funnels. To overcome perceptions of ML/RL model being black boxes, marketers are enabled to customize ML/RL optimization by defining the states of the funnel and the actions corresponding to each state. Also, marketers can control the optimization, e.g., by ensuring that ML/RL models do not define new states, change marketer defined states or replace defined actions without marketer approval.

ML/RL model training from the scratch can suffer from a cold-start problem. For instance, at the beginning of a campaign, the actions predicted by the model are associated with low confidence, e.g., due to insufficient volume of training data. The present techniques solve the cold -start problem by relying on the conversion funnel states, actions, and associated costs to be initially defined by marketers at the beginning of an ad campaign, when there is limited or no training data. As the model trains on more data, the probability of executing optimal actions predicted by the model increases and reliance on marketer initially defined actions decreases.

This design also reduces the need to train the ML/RL algorithms with huge initial volumes of data, and therefore marketers can manage transition from defined actions to actions predicted using machine learning models.

Training the described ML/RL model to predict actions at each state along with cost of each action is comparable to training a pCTR (predict click through rate) model via multi-layer perceptron neural networks and requires significantly less initial data than traditional ML/RL models. Due to reduced initial training data requirements and faster training cycles, the described techniques can be applied to small and medium sized campaigns across diverse industries, products, and regions.

The present techniques for optimizing the conversion funnel are easy to understand by marketing professionals, guided by human (i.e., marketer) expertise, and perform better than direct application of machine learning models without incorporating the human expertise. Marketer expertise is particularly helpful at the beginning of an ad campaign, when the volume of available training data is limited.

Companies that provide advertising technology can use these techniques to define a new type of campaign, e.g., a multi-stage conversion campaign, or advertising product, to directly support the conversion funnel. Offline advertising can also use these techniques when the relevant signals are available online in near real time. For example, such signals can include, with user permission, user visits to a store, user purchases of certain items, etc.

Further to the descriptions above, a user may be provided with controls allowing the user to make an election as to both if and when systems, programs or features described herein may enable collection of user information (e.g., information about a user's social network, social actions or activities, profession, a user's preferences, or a user's current location), and if the user is sent content or communications from a server. In addition, certain data may be treated in one or more ways before it is stored or used, so that personally identifiable information is removed. For example, a user's identity may be treated so that no personally identifiable information can be determined for the user, or a user's geographic location may be generalized where location information is obtained (such as to a city, ZIP code, or state level), so that a particular location of a user cannot be determined. Thus, the user may have control over what information is collected about the user, how that information is used, and what information is provided to the user.

CONCLUSION

Techniques described incorporate initial inputs from marketers and apply machine learning techniques to optimize conversion funnels in the context of online advertising. Marketers initially apply insights into the process of optimizing actions at each state of the funnel and start by manually customizing the optimization, e.g., by defining the states and optimal actions to move consumers to next states of a funnel. With user permission, a machine learning model that uses Machine Learning/Reinforced learning is randomly applied to try other actions defined by marketers (e.g., for other states) to collect training data on various actions. The training data auto collected is used to train the model and eliminates the requirement of a large amount of training data to be provided at the beginning of the process. The trained model thus obtained predicts optimal actions to move the consumer through the conversion funnel. Predicted actions using the model for a given state, along with supporting evidence, are provided to marketers for review and approval. Marketers can modify the funnel parameters, e.g., states and actions, based on such evidence.