Chapter 3
Reinforcement Learning in Social Media Marketing

Patrik Eklund
https://orcid.org/0000-0003-3965-2834
Umea University, Sweden

ABSTRACT

In this chapter, the authors describe an architecture for reinforcement learning in social media marketing. The rule bases used for action selection within the architecture build upon many-valued (fuzzy) logic. Action evaluation and internal learning is based on neural network like structures. In using variables measuring the effect of advertising, we must understand direction of influence between advertiser, owning the content of the advertisement, and advertisee, as the target of an advertisement, and as facilitated by social media marketing. Examples are drawn from Facebook marketing.

INTRODUCTION

Social media marketing (SMM) is a special form of marketing using combinations of various media channels. General marketing AI principles are relevant, but the underlying information structures need to be specific to SMM, and are further specified within particular SMM media e.g. like in Facebook.

Given advertiser objectives and expected outcomes, including return of investment (ROI) expectations, the advertiser specifies the overall campaign budget, and other investments related to the execution of the marketing campaign.

Various constraints are provided, many of which fall outside the scope of the way the campaign is handled and executed.

In the end, ROI depends on conversion, even if conversions alone are not sufficient to calculate the ROI of a campaign. Further, ROI can be seen as static and global with respect to the campaign as a whole, or as dynamic and local given conversion frequency appearing during a campaign.

An important part of advertiser objectives relates to customer relations and its management (CRM), which embraces management of social media marketing audiences. Audience and CRM modelling is continuously on the agenda among marketing technology providers, but underlying structure, phenom-

DOI: 10.4018/978-1-7998-5077-9.ch003

ena and dynamics is not always fully understood. This leads to approaches like those involving social media marketing campaigns, where a recorded dialogue between campaign and audience improves the understanding of the relation between the advertiser (or its brand), and the advertisee.

The objectives of this chapter is to show how reinforcement learning involving fuzzy logic rule bases can provide SMM ROI optimization.

BACKGROUND

An Ad conveys a message, and is represented by keywords and the content of the Ad, containing text, images and video, annotated with links and objects (like buttons) enabling and inviting the Advertisee to interact with the Ad, and become influenced by the Ad. An Ad is like an expression presented (syntactically), and perceived (semantically) by the Advertisee. That perception depends on the Advertisee profile and behaviour. Ads with video content are either digital video Ads or as based on programmatic video advertising.

The content of that Ad can be dynamically modified, and the characteristics of the Ad should be described as fine-granularly as possible. A coarse-granular description of the Ad leaves less room for subjecting the Ad itself to modification during optimization of the campaign.

Ads are part of Ad sets, and a set of Ad sets constitutes a Campaign. This means that Advertisers must deal not just with the fine-granular description of respective Ads, but also consider the characteristics of the Ad sets, and the structure of the Campaign.

Facebook generates cost related to Ads, so that costs of Ad sets are aggregated costs. A campaign cost is the accumulation of all costs incurred by specific Ads and Ad sets.

Bidding and Facebook's Ads auction is not only price related, i.e., the highest monetary bid does not necessarily win, but the winner is the Ad that creates the most overall value. The mechanism for calculating that "overall value" is hidden. Apart from the Advertiser Bid, being manual or automatic, Ad quality and relevance comes into play. An auction takes place whenever someone is eligible to see an Ad. Manual bidding influences reinforcement learning differently as compared to automatic bidding.

Ad performance can be measured e.g. by engagement and brand lift. Engagement with the Ad, can be measured either by hovering, scrolling or interaction with an extension, scroll velocity as a proxy for attention, how long a native Ad is viewed (even if not clicked), average time reading or watching, and so on. Brand lift can be measure e.g. by shares, followings, and email subscriptions.

There is a continuously on-going debate about performance metrics, but the *What* of measuring seems not properly connected with the *Why* of measuring. Metrics has been focused on outcome and performance, but the trend is to focus more also on Ad content. Formal terminology and nomenclatures, and as based on formal logic [Eklund et al 2014], for classifying Ad content is largely missing in marketing.

Performance metrics include

- Actions
- Clicks
- CTR
- Impressions
- Relevance score