

Determining optimal advertisement frequency capping policy via Markov decision processes to maximize click through rates

James G. Shanahan
Independent Consultant (USA)

Dirk Van den Poel
Ghent University (Belgium)

Feature2	D2; \$50	D3; \$350
	Policy = D1 Rev= \$100	
D4; \$550		
D5; \$1550		
Feature1		

Introduction

The term "**frequency capping**" refers to restricting (capping) the amount of times (frequency) a specific visitor is shown a particular advertisement within a period of time. For example, a frequency cap of "**3 per 24**" for an ad means that after exposing the user to the same ad 3 times, the visitor will not be shown that ad for 24 hours.

Some of the key benefits of frequency capping are:

- **Increased reach** – frequency capping is an efficient way to increase reach by expanding the number of unique users that see the ad.
- **Prevent burnout**– Click Through Rates (CTR) have been shown to drop precipitously after the first exposure and to plateau at four because of "ad blindness", wherein the user becomes so used to seeing the ad that the user essentially becomes blind to that ad.
- **Maintain quality scores**–maintain a competitive quality score, a core component of expected CPM-based ranking.

Currently, the frequency capping policy for an ad is **heuristically** set and is optimized for **short term gain** and set globally (one-size-fits-all).

Propose to set frequency capping policies for different online marketing segments using **Markov decision processes** (MDP). Advantages:

- A data driven approach that optimizes the **life time value** of site visitors.
- Optimal policy for each segment

Methodology

Goal: Maximize long term revenue (CTR) of site visitor by treating problem as a joint optimization of both visitor segmentation and frequency cap policy determination (similar to [1] for direct mail).




Figure 1. Search space of possible segmentations and select segmentation that provides highest long term value as determined by MDP. This solution corresponds to a frequency cap policy for each segment.

Each visitor is represented in term of variables such as (RFM data):

- **Recency (R)** – Number of impressions of this ad without a click by user
- **Frequency (F)** – How many times has this user been served ads from the topic of the ad before
- **Monetary (M)** – How many times did this user click on any ad? Revenue?
- **Others:** How long do I know this user? Average dwell time? How many times has this user been served before? Performance on categories (ad, pub). Demographic etc.

Every visitor in a segment is assumed to be in the same state. People migrate between segments, i.e. people move from one state to another based on behavior (e.g., click). Observe the states at given time intervals. Decide on the frequency of ad impressions in the next time interval.

Operationalization

Use genetic algorithms to search through possible segmentations of visitors (k-Means, Dtree) where the fitness of each candidate solution is the total discounted revenue over an infinite time horizon.

GA-MDP ALGORITHM

Input: RFM dataset

Output: Segmentation of visitors (variables and breakpoints), and policy for each segment

- Create initial pool of candidate segmentations (size = 100)
- REPEAT
 - For each candidate segmentation
 - Set up data
 1. Draw 100 bootstrap samples
 2. Transition prob. matrices P_{ij}^a ,
 3. Reward vector R_i^a
 - Determine optimal policy via MDP
 - fitness =total discounted revenue (mean – stdev)
 - Create next generation using mutation and crossover
- UNTIL set number of generations

MDP (S, D, T, R, $\gamma=0.9$, Horizon=infinite)

S: possible state of visitor at time t

D: Frequency capping policy during the next time period given the state observed at the end of this time period.

P: Probability to move from state S_t to state S_{t+1} given the frequency capping policy

R: Observed revenue (RPM) from the visitor in each state when policy is applied in the current period

VALUE ITERATION (repeat til convergence):

$$V^n(s) = \begin{cases} \max_{a \in A} \left\{ R(s, a) + \gamma \sum_{s' \in S} P(s'|s, a) V^{n-1}(s') \right\} & n > 0 \\ 0 & n = 0 \end{cases}$$

OPTIMAL POLICY

$$D^*(s) = \arg \max_{a \in A} \left\{ R(s, a) + \gamma \sum_{s' \in S} P(s'|s, a) V^{n-1}(s') \right\}$$

Experiments

Experiments on **frequency capping** work is ongoing at major publisher.©

Loyalty card for apparel retailer. Policies:

- (0) Do nothing;
- (1) Send thank you letter
- (2) Send letter plus 1 movie tickets
- (3) Send letter plus 3 movie tickets
- (4) Send thank you letter plus expensive dinner.

Eight variables: length of relationship, recency of last purchase, total purchased (over 3 time periods), frequency of purchases. (13,000 customers)

Table 1. Optimal segmentation (just over lifetime spend variable) and corresponding optimal policy.

Visitor Segment	1	2	3	4	5
Centroid	45.72	228.16	477.34	863.96	1668.42
Policy	1	2	3	1	4
Size	6743	3848	1944	779	136
Value function	2753.39	2990.38	3328.31	3571.42	3815.10
Profit	18566109	11506982	6470235	2782136	518854

Conclusions

Setting frequency capping policies using the proposed approach optimizes the long term revenue of a site visitor.

The experiments on frequency capping are ongoing. Results on loyalty card program are very encouraging.

Literature

- [1] J.J. Jonker, N. Piersma, and D. Van den Poel, "Joint Optimization of Customer Segmentation and Marketing Policy to Maximize Long-Term Profitability," *Expert Systems with Applications*, vol. 27, 2004, pp. 159-168.

For further information

Please contact James.Shanahan@gmail.com;
Dirk.VandenPoel@UGent.be.