World Academy of Science, Engineering and Technology
International Journal of Economics and Management Engineering
Vol:17, No:3, 2023

Personalized Email Marketing Strategy: A Reinforcement Learning Approach

Lei Zhang, Tingting Xu, Jun He, Zhenyu Yan, Roger Brooks

Abstract**—**Email marketing is one of the most important segments of online marketing. Email content is vital to customers. Different customers may have different familiarity with a product, so a successful marketing strategy must personalize email content based on individual customers' product affinity. In this study, we build our personalized email marketing strategy with three types of emails: nurture, promotion, and conversion. Each type of emails has a different influence on customers. We investigate this difference by analyzing customers' open rates, click rates and opt-out rates. Feature importance from response models is also analyzed. The goal of the marketing strategy is to improve the click rate on conversion-type emails. To build the personalized strategy, we formulate the problem as a reinforcement learning problem and adopt a Q-learning algorithm with variations. The simulation results show that our model-based strategy outperforms the current marketer's strategy.

Keywords**—**Email marketing, email content, reinforcement learning, machine learning, Q-learning

I. Introduction

OVER the past decade, online digital media advertising grows significantly and gradually replaces the traditional media advertising. Online digital media advertising allows marketers to communicate with their customers more directly and effectively [1]. Marketers are able to reach their customers via different digital channels, such as webpages, email, YouTube. With the rise of social channels like Facebook and Twitter where marketers and customers can interact with each other, it is much easier for marketers to target their customers, particularly young customers, comparing with traditional advertising channels [2].

Among all of the digital channels, email channel is one of the most widely used channels. Compared with printed mail, the cost of sending a large amount of email is considered to be marginal [3]. It is estimated that nearly 4 billion people in the world have their own email accounts and over 293 billion emails were sent or received per day in 2019 [4]. Email marketing has also been proven to be more effective than other channels in connecting business with customers. It has at least twice the return on investment of other main forms of online marketing methods such as webpage banners, search engine marketing, keyword advertising, etc. [5]

In this study, we focus on email marketing for the Adobe Creative Cloud. Adobe Creative Could email marketing mainly supports Adobe's designer products, such as Photoshop, Illustrator, Premiere, etc. [6] The goal of the email campaign is to maintain and increase product awareness and sales. In the email campaign, there are multiple types of emails designed for different purposes as in Fig. 1.

Nurture: Nurture-type email helps customers learn the products and services. It educates the customers as to why they should buy or try this product. The nurture-type emails typically include the latest product updates and features, feedback from other customers, and tutorials on new features.

Promotion: Promotion-type emails allow customers to explore the full features of the product at a heavily discounted price. During the promotion period, customers have the opportunity to try the premium features and determine if they are worth buying. Promotion-type emails help convert customers from free users to paid subscribers by increasing the attractiveness of products through discounts.

Conversion: Conversion-type emails attempt to convert free customers to the paid customers. It is the ultimate goal of the email campaign. Although sending conversion-type emails to customers has the potential to generate the most value, it is also the riskiest. For a new customer who is just starting to learn the product, sending too many conversion-type emails may not get them to convert. Instead, they may feel abrupt, lose interest in learning the product, and eventually opt out of the email list. This is also verified in the analysis in a later section.

Sending emails is free, and sending emails can be abused. Bombing too many emails within a short period of time could affect customers' experience and make them opt out from the list. According to a recent study [7], the top reason for email opt-outs is excessive frequency, accounting for 46.4% of total opt-outs. Marketers should find the right contact frequency for each recipient to avoid putting too much pressure on them.

This study focuses on optimizing the email delivery policy for each customer. The goal is to improve the click rate of the conversion-type emails, since customers who click on such emails have a high chance of converting from free to paid customers at a later stage. We applied a reinforcement learning approach to tackle the problem. Reinforcement learning is a machine learning method used to learn how to make the best action to maximize the rewards over a time horizon [8]. We adopt Q-learning algorithms with the concept of experience reply to train models based on historical marketer's experience and customers' responses [9]. We also applied the latest variation of Q-learning algorithm to overcome the over-optimism issue.

Lei Zhang, Tingting Xu, Jun He, and Zhenyu Yan are with the Department of Media & Advertising Solutions, Adobe Inc, San Jose, CA, 95110, USA (e-mail: lzhang@adobe.com, tixu@adobe.com, juhe@adobe.com, wyan@adobe.com, robrooks@adobe.com).

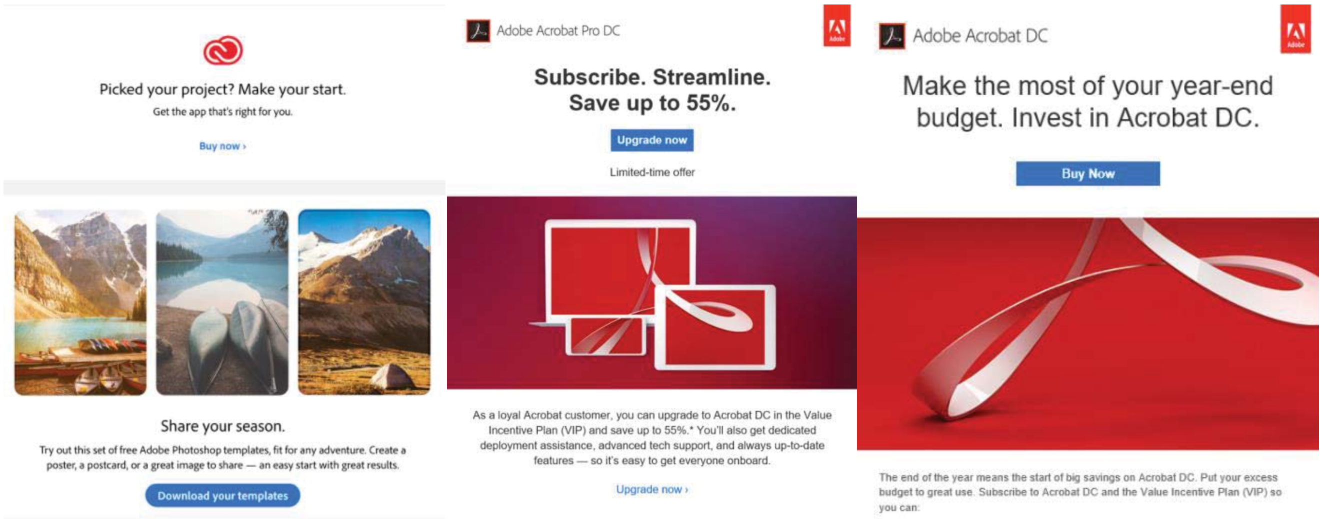

Fig. 1 Sample emails of: nurture (a), promotion (b), conversion (c) types

II. Related Work

In recent years, email marketing has been widely studied in different aspects. Chittenden and Rettie researched factors affecting recipients' response in [10]. They summarized three stages of effective email marketing: getting recipients to open the e-mail, keeping their interest and convincing them to respond.

Gupta et al. investigated the optimal amount of email to customers in [11]. They analyzed the cost and benefits of sending each type of email and formulated the problem as A multi-object optimization to find the optimal sending volume.

Yang et al. investigated the important features of post-stay emails for hotels in [12]. They find that standardized and system-generated emails were difficult to impress guests, but personalized emails had the strongest effect on attitude towards a hotel brand and revisit intention. They encourage that companies should create more targeted and personalized emails, which can build long-term customer loyalty.

Jaidka et al. studied the differences in recipients' preferences for subject lines of marketing emails across different industries, based on the idea that different styles of subject lines may have different click-through rate [13]. They proposed a language model to predict the click-through rate given a specific email subject line in a business domain. The language model can also be used to detect clickbait articles.

Litinthong et al. investigated the impact of email marketing on online purchase behavior [14]. They concluded that the three dimensions of email marketing, namely newsletters, promotion, and viral marketing, are all positively correlated with online purchase behavior. Moshe et al. studied the effect of color on recipients' response to emails [15]. They suggest that appropriate colors in emails are important to capture the recipient's attention and create the feeling that induces a positive response to requests and marketing communications.

Micheaux demonstrated how email contact frequency affects the customer response through a set of specifically designed experiments [16]. It is found that the relevance of the email content is more important than contact frequency. High contact frequency would not result in unsubscribes if the content is relevant to the recipient's interest. It is suggested that email marketers should deliver more targeted and personalized content, rather than tuning contact frequency.

While there are many separate studies on the effectiveness of content and frequency of email marketing, very few work has been done on generating a comprehensive strategy that can personalize email content and contact frequency to promote user conversion. This is not unexpected. As the scope of the problem becomes broader, the complicity of the problem increases exponentially.

We propose a framework based on reinforcement learning that learns from the sequences of historical delivered emails and recipients' responses. By fully training with these sequences of actions and responses, the model can generate personalized email delivery decisions that are the best for a given recipient, taking into account whether the recipient should be contacted and which content type of email should be delivered.

III. Data Insight

To understand how recipients respond to different content type of emails, the data insight is provided in this section. We analyzed the open rates, click rates and opt-out rates of different email types, and also built a random forest model to find important features that could contribute to the clicks on conversion-type emails.

TABLE I
Email Response Rates of Different Contents

Email Type	Sent	Open Rate	Click Rate	Opt-out Rate
Nurture	500,135	35.82%	0.31%	0.05%
Promotion	862,090	26.62%	0.29%	0.01%
Conversion	790,252	32.10%	0.30%	0.17%

We tracked a total of 179,484 users for 50 days and collected the deliveries and user responses of over 2 million emails, as shown in Table I. It can be seen that nurture-type emails have the highest open rate and a relatively low opt-out rate. This is understandable, as nurture-type content brings up-to-date product information and usage tips to customers, making it easier for customers to use the product. It is also less likely to cause stress to the customers.

Emails with promotion-type content has the lowest open rate and opt-out rate. This result is consistent with Micheaux's study [16]. This is because promotional emails are targeted only to those users with a specific demand. Those users who find the promotional content irrelevant will simply ignore the email, which leads to low open rates. It also has the lowest opt-out rate because promotional emails rarely irritate users, even to those users who find the content irrelevant.

Conversion-type emails convert free or trial users to paid subscribers. This is the most valuable type of email for marketers. However, it also puts a lot of pressure on recipients who are not prepared to pay for the product. It has an opt-out rate over 3 times higher than nurture-type emails and 17 times higher than promotion-type emails.

For all three types of emails, the click rates of emails are very close. A click is a further action taken by the recipient after opening and viewing the email content. To better understand how the recipients like the content after opening an email, we calculate the click-through rate by dividing click rate by the open rate. We also calculate post-open opt-out rate email by dividing opt-out rate by the open rate.

TABLE II
EMAIL CLICK-THROUGH RATES OF DIFFERENT CONTENTS

Email Type	Click-through Rate	Post-open Opt-out Rate
Nurture	0.87%	0.140%
Promotion	1.09%	0.038%
Conversion	0.93%	0.530%

From Table II, we can see that the content of promotional emails has the most positive impression on the recipients, as it has highest click-through rate and lowest post-open opt-out rate. For nurture-type emails, the click-through rate is moderate and the post-open opt-out rate is low because the recipients can expect the content when opening nurture-type emails.

The conversion-type email is controversial because it has both a high click-through rate and a high post-open opt-out rate. This means that for different recipients, their attitude towards the same conversion-type email can be different. For some recipients, the content creates positive feelings, while for others, the same content only creates stress.

As pointed out in [16], "marketers should change the content with the objective of increasing advertising relevance." We think the relevance, i.e., the customer's attitude towards the product, is the key to the conversion-type emails. Marketers should carefully consider whether the conversion-type emails are relevant to the recipient before sending those emails, as their potentially high pressure could lead to the permanent loss of potential customers.

Based on the above analysis, the key to building a successful marketing campaign is to understand the relevance of the conversion-type emails to recipients. Since high relevance often leads to high click-through rates, the problem can be converted to understand what factors can drive the click-through rates for conversion-type emails.

We collected a set of records of recipients who opened the conversion-type emails and generated a set of features about recipients' previous responses to different types of emails, such as the number of opens on nurture-type emails (nOpen_N), numbers of clicks on promotion-type emails (nClick_P), …etc. (see Table III). We then trained a logistic regression model to predict whether the conversion-type email is clicked or not using these features as variables and analyzed the variable importance to conversion-type email clicks.

$$Prob_{click} \sim \{nOpen\_N, nOpen\_P, nOpen\_C, \dots\} \qquad (1)$$

The variable importance result is shown in Fig. 2. We can see that the top 3 important variables are all number of recipient's historical clicks on different type of emails. The number of clicks on promotion-type emails and nurture-type emails is even more important than the number of clicks on conversion-type emails. This finding suggests that the click probability of conversion-type email is highly correlated with the number of clicks on historical nurture-type and promotion-type emails. It also suggests that in order to improve the relevance between recipients and conversion-type emails, we can increase the recipient's clicks on nurture-type and promotion-type emails by sending more of these types of emails, as they are more likely to be accepted by recipients with lower relevance.

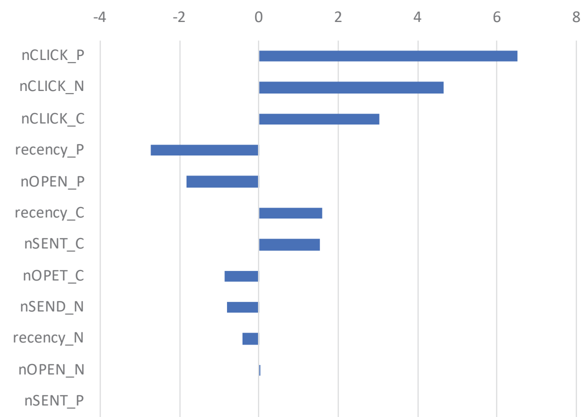

Fig. 2 Most Important Features for Conversion Click

TABLE III
FEATURE LISTS

Feature name	Description
nOpen	The recipient's total opens in the past
nClick	The recipient's total clicks in the past
nSent	Total emails sent to that recipient in the past
recency	The days between the last two emails sent
suffix '_N', '_P', '_C'	_N: Nurture-type email; '_P' Promotion-type email; '_C' Conversion-type email

IV. REINFORCEMENT LEARNING ARCHITECTURE

In this section, we demonstrate how we formulate the personalized email delivery problem using a reinforcement learning framework, and how we solve the problem.

A. Reinforcement Learning and Q-Learning

Reinforcement learning is a set of machine learning algorithms that learn from the historical experience. Unlike supervised machine learning algorithms that require manual labeling of positive and negative samples, reinforcement learning algorithms formulate the problem as a trial-and-error process where agents adjust their actions based on feedback rewards associated with historical actions. [8]

The goal of the reinforcement learning is to find the optimal action polity $\pi = \pi(a|s)$ that could maximize the expected total rewards over a time horizon:

$$\pi(s) = arg \max_{\pi \in \prod} \mathbb{E}[\textstyle\sum_{t=0}^{\infty} \gamma^t r_t] \tag{2}$$

where s: state variable; γ: discount factor, trade-off between instantaneous and future rewards; r_t: instantaneous reward at time t; π: action policy or rule to determines actions given states; $\mathbb{E}$ [·]: the expectation over randomness of state transition.

Q-learning [17] is an off policy, temporal difference (TD) reinforcement learning algorithm which learns from the historical actions and rewards. Q-learning uses a Q function to approximate the reward given a sequence of actions, which can be represented as:

$$Q^{\pi}(s,a) = \mathbb{E}[\textstyle\sum_{t=0}^{\infty} \gamma^t r_t \,|s_0 = s, a_0 = a] \tag{3}$$

The initial $Q^{\pi}(s,a)$ can be trained from historically data, by solving bellman equation:

$$Q^{\pi}(s,a) = r(s,a) + \gamma \textstyle\sum_{s' \in S} P(s'|s,a) Q^{\pi}\big(s', \pi(s')\big) \tag{4}$$

where $P(s'|s,a)$ is the transition probability from the state s to state s' after taking action a.

The optimal policy also obeys the dynamic programming principle:

$$Q^*(s,a) = r(s,a) + \gamma \textstyle\sum_{s' \in S} P(s'|s,a) \max_{a \in A} Q^*(s',a) \tag{5}$$

with the $Q^*(s,a)$, given the current state s_t, the best action a_t^* can be find as:

$$a_t^* = arg\max_{a \in \mathrm{A}} Q^*(s,a) \tag{6}$$

B. Problem Formulation

We formulate the email marketing problem using the reinforcement learning framework. The three key elements of reinforcement learning are state, action and rewards. For each recipient, the state variable consists of recipient's profile and historical contacts and responses information as described in Table IV.

To a marketer, the actions are the ways to make contact with their customers. A marketer can choose to give up any opportunity with a customer if he/she feels the contact frequency is relatively high to that customer.

TABLE IV
STATE VARIABLE

Feature name	Description
Age	The recipient's age
Sex	Male or Female
…….	……..
nOpen, _N, _P, _C	The recipient's total open in the past, for different email types
nClick, _N, _P, _C	The recipient's total clicks in the past, for different email types
nSent, _N, _P, _C	Total emails sent to that recipient in the past, for different email types
recency, _N, _P, _C	The days between the last two emails sent
….	_......

TABLE V
ACTION LIST

Action	Description
Send nurture-type email	Send nurture-type emails to the current recipient
Send promotion-type email	Send promotion-type emails to the current recipient
Send conversion-type email	Send conversion-type emails to the current recipient
Not Send	Do not send any email, at the current decision time

The design of reward function is usually a challenging part in reinforcement learning algorithms. Reward functions describe the incentive or punishment for an agent's actions. The reward function will guide the agent's actions in the preferred direction.

In the considered email marketing problem, the ultimate goal is to drive recipients to click on the conversion-type emails. This will give the marketer the highest rewards. However, as described in Section II, conversion-type emails may irritate some recipients who are not ready yet. Therefore, marketers should avoid sending conversion-type emails to these customers. Any opt-out recipients will cause marketers to receive a huge penalty. Besides, recipients' clicks on nurture-type and promotion-type emails are also important, which indirectly drive conversion-type email click-through rates. These responses are also rewarded with relatively high positive returns. We also consider recipients' open responses on any type of emails as positive feedback, because they are a prerequisite for click behavior. All open actions are rewarded with small positive returns. We also expect marketers to control their sending frequency so that each delivery is reasonable, so we add a small penalty for sending any type of emails. The relative rewards of all actions are shown in Table VI.

TABLE VI
REWARD TABLE

Reward	Nurture	Promotion	Conversion
Open	small reward	small reward	small reward
Click	moderate reward	moderate reward	high reward
Opt-out	high penalty	high penalty	high penalty
Send email	small penalty	small penalty	small penalty

C. Training the Q-Network

The algorithm is based on the optimal Q-function $Q^*(s,a)$, which maximizes the return that can be obtained by taking action a at state s and following the optimal policy thereafter.

The optimal Q function $Q^*(s,a)$ can be approximated using a neural network. However, it is difficult to solve the TD equation (5) directly since the $Q^*(s,a)$ exist on both sides of the equation, and training directly will make the result unstable. In [18], Mnih et al. proposed to adopt two neural networks for training and target networks separately to solve the Q function learning task. Furthermore, training Q-network with the above idea tends to select overestimated values, resulting in over-optimistic value estimates. To overcome this, we adopt the double Q-learning method proposed by Hasselt et al. [19], which decouples the action selection from the evaluation:

$$Y_t = r_t + \gamma Q_{TargetNet}\left(S_{t+1}, a_{t+1} = arg\max_{a\in A} Q(s_{t+1}, a; \theta_t); \theta_t^-\right) \quad (7)$$

For the network structure, we tried both single-layer and double-layer neural networks. Following the idea of experience replay, we formulated the training data as $< s_t, a_t, s_{t+1}, r_t >$ tuple, so that only these variables are involved in the training process. Fig. 3 shows the training loss in a neural network training process.

Fig. 3 Neural Network Training Process

V. Simulation Result

In this section, we will show the simulation workflow. We firstly build a recipient response simulator to estimate recipients' responses. Next, we run the simulator with actions learnt from Reinforcement Learning models and collect rewards based on predicted recipients' responses. At last, we will compare the rewards of proposed method, estimated actual marketer's policy, and marketer's actual actions.

A. Recipient Simulator

Recipient simulator is a group of models that estimate recipients' responses. There are three predictive models in the recipient simulator, open model, click model and opt-out model. Open model predicts the probability of the recipient opening this email given the current state s_t and the email type (nurture, promotion or conversion). The click model and opt-out model predict the probability of the recipient clicking and opting-out the email given that the recipient opens the email.

The output of all the predictive models is probabilities, in order to get a binary value to indicate whether the response is positive or negative, we set a random binary variable sampling process where the probability of generating positive response comes from the predictive models. The three models work in a sequential pattern, described in Fig. 4.

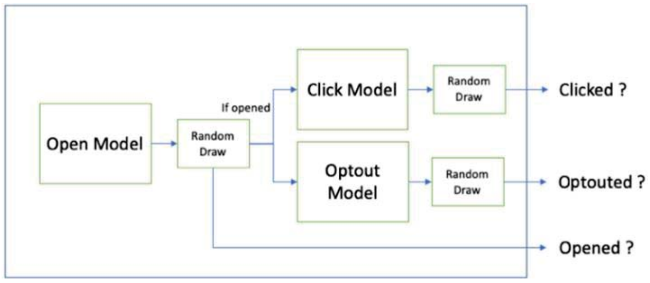

Fig. 4 The Recipient Simulator

We build all the models with logistic regression (LR) and random forest (RF). The AUC of all models are around 0.75. We consider that this accuracy is decent for the simulator, as seeking high accuracy is not the focus of this paper.

B. Simulation Workflow

The email marketing delivery simulation workflow is shown in Fig. 5. Given a recipient with historical data, the model helps the marketer decide actions how to send emails at the current time. Recipient receives the email and generates the response according to the response simulator. The marketer receives the response, collect rewards and makes the next sending decision.

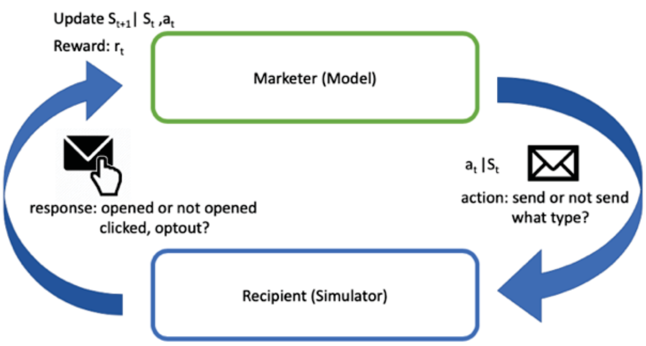

Fig. 5 The simulation workflow

C. Simulation Result

To test the proposed method, we prepared a dataset with 179,484 recipients, different initial states according to their historical data. We run simulation for 50 days. For each day, marketer can determine actions based on different policies or rules, and we will compare the following three policies:

- *Actual policy*: is to use real marketer's policy in the dataset and real recipient's responses to calculate the result.
- *Base policy*: is to use real marketer's policy in the dataset but use simulator's response to calculate the result. Since the simulator may have its own prediction error, the base policy shows the real marketer's performance, with

elimination of simulator’s error, when it is used to compare with the model policy.

- *RL policy*: is to use model’s policy and simulator to calculate the result.

For fair competition, we also tuned the model parameters so that the model’s sent frequency is close to the real marketer’s sent frequency. Table VII shows the email sent percentage of different email types, over the 50 days.

TABLE VII
EMAIL SENT PERCENTAGE

	Unsend	Nurture	Promotion	Conversion
Real Marketer	75%	5.87%	10.11%	9.27%
RL-Policy	78%	10.08%	3.59%	8.43%

The evaluation result is shown in Fig. 6. Compared with the base policy, the open rate of all types of emails are improved using the RL policy. The improvement on conversion-type emails is the most significant. We consider this is related to the reward setting, as the objective is to drive clicks on conversion-type emails. Meanwhile, we can also see that the open rates of the base policy and actual policy are close. This suggests that the response simulator has similar behaviors to the real actual recipients.

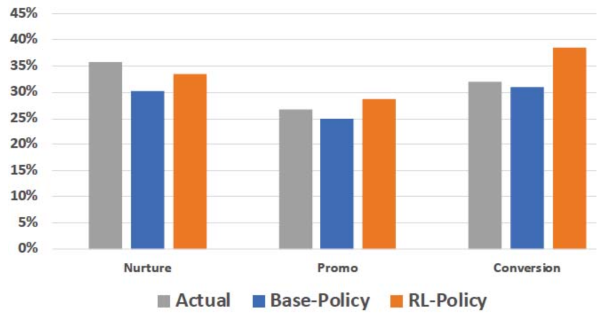

Fig. 6 Open rates of different policies

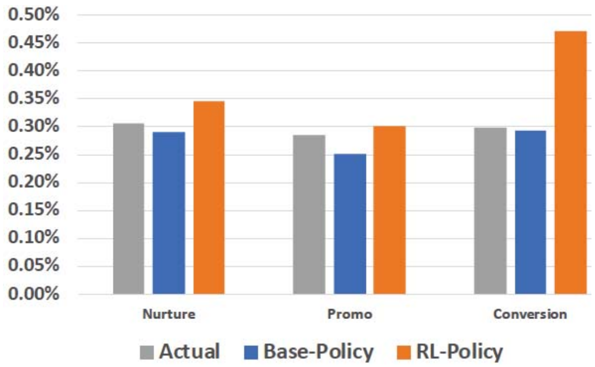

Fig. 7 Click rates of different policies

Fig. 7 shows the click rates of different policies. The RL policy outperforms the alternatives on all three types of emails. The click rate of conversion-type emails is largely improved from 0.30% to 0.45%. We consider the model works strategically to drive recipients to click on conversion-type emails in a statistic view. It is also worth to mention that the performance of actual policy and base policy is quite close, which means the simulator does not overestimate the recipients’ responses.

Fig. 8 shows the distribution of recipients’ conversion-type email click rate. With the actual policy, the percentage of recipients with non-zero conversion clicks is 1.27%, and using the RL-policy, the percentage of recipients with non-zero conversion click is increased to 1.55%. This suggests that RL policy can not only improve conversion clicks, but also convert recipients with zero click before.

To further evaluate our model, we run a different test. We used the same recipients as before, but randomly picked days on their history as the initial days. Therefore, the dataset is more unbalanced and randomized. We still evaluate the performance on actual policy, base-policy, and RL-policy. Figs. 9 and 10 show the open rates and click rates of different policy.

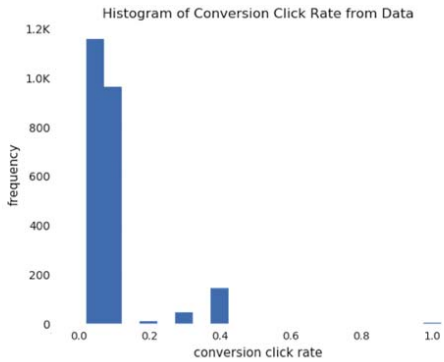

(a) Actual policy

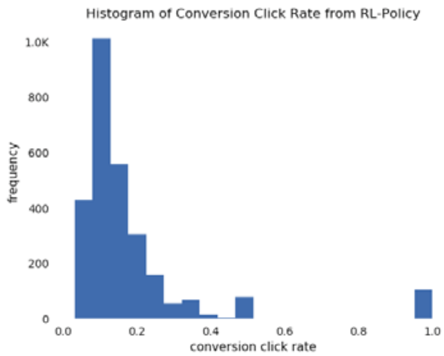

(b) RL policy

Fig. 8 The distribution of recipient’s click rate

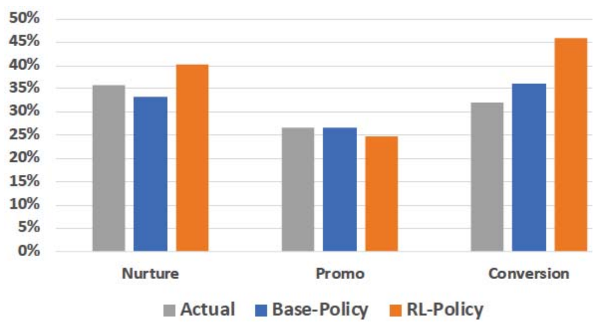

Fig. 9 Open rates of different policies on random starting days

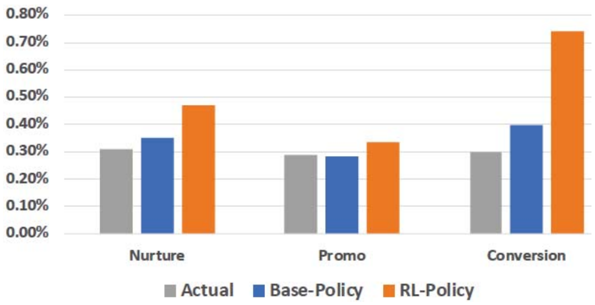

Fig. 10 Click rates of different policies on random starting days

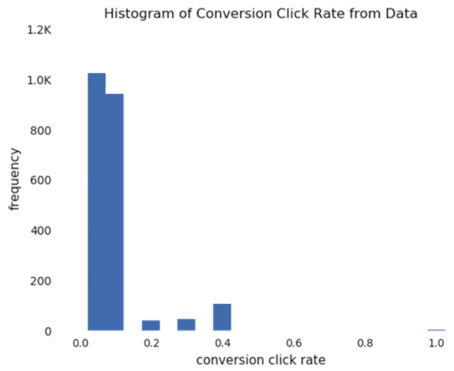

(a) Actual policy

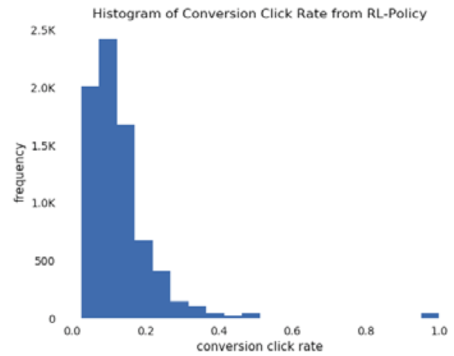

(b) RL policy

Fig. 11 The Distribution of recipient's conversion-type email click rate on random starting days

We can find that the improvement of RL-policy is larger than all recipients starting on the same day. We consider the reason behind this is the model can generate more personalized email delivery policy than the real marketer. The sooner the model engaged, the better result can be obtained. This conclusion can also be supported by checking distribution of recipients' conversion clicks, shown in Fig. 11. It shows the distribution of conversion click rate of both base policy and RL-policy. The percentage of recipients with non-zero conversion clicks is 4.14% with RL-policy and is 1.27% with actual policy. RL-policy has more than 3 times improvement on first conversion click recipients, showing its ability of converting potential paid users.

Fig. 12 shows a segment of mail delivery policy the proposed model generated for a recipient. During the first several contacts, the model sent nurture-type and promotion-type emails. For nurture-type emails, there are positive responses, i.e., the recipient opened the nurtures emails. Then, the model sent a conversion-type email, but there is no response from the recipient. The model switched back and sent more nurture-type emails, and the recipient clicked the email. Having observed this, the model sent conversion-type emails, and the recipient opened one of them. Then model alternatingly sent nurture and conversion-type emails and finally the recipient clicked the conversion-type email.

The model shows several sending patterns: first, the model does not send conversion-type emails until there are enough positive responses, either opens or clicks from the recipients. Second, the model sends conversion-type emails and other emails alternatingly to prevent opt-out by keeping low pressure.

VI. Conclusion and Future Work

In this study, we investigate the influence of different email types on customers: nurture, promotion and conversion. Our analysis suggests that the attitude of recipients is a key factor in motivating them to click on conversion emails. In order to positively influence recipients' attitudes, we recommend implementing a strategy of sending both nurture and promotion emails. This approach allows us to not only enhance customers' knowledge and experience of the product, but also create a less intrusive and more engaging email experience. We then formulate the email marketing problem into a reinforcement learning framework and applied Q-learning algorithm to build a marketing model based on our findings. The simulation results show that our designed model's delivery policy outperforms that of a human marketer. To further validate our findings, we plan to deploy the model in a real-world marketing channel and evaluate its performance.

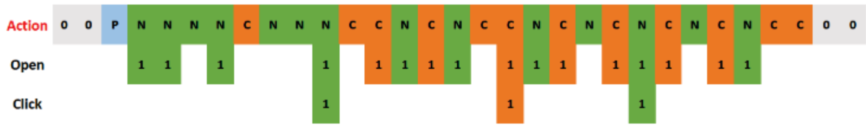

Fig. 12 A Segment of continuous model policy (N: send nurture-type email; P: send promotion-type email; C: send conversion-type email; 0: not send)

References

[1] R. Ahmed, D. Streimikiene, G. Berchtold, J. Vveinhardt, R. Soomro and Z. Channar, "Effectiveness of Online Digital Media Advertising as A Strategic Tool for Building Brand Sustainability: Evidence from FMCGs and Services Sectors of Pakistan", Sustainability, vol. 11, pp. 1-40, 2019.

[2] L. Fuxman, H. Elifoglu, C. Chao and T. Li, "Digital Advertising: A More Effective Way to Promote Businesses' Products", Journal of Business

Administration Research, vol. 3, No. 2, pp. 59 – 67, 2014
[3] M. Merisavo, M. Raulas, "The impact of email marketing on brand loyalty", Journal of Product and Brand Management, vol 13, No. 7, pp.498-505, 2004.
[4] The Radicati Group, Inc. "Email Statistics Report, 2019-2023", https://www.radicati.com/wp/wp-content/uploads/2018/12/Email-Statistics-Report-2019-2023-Executive-Summary.pdf
[5] C.H., Cho, H.K., Khang, "The State of Internet-Related Research in Communications, Marketing, and Advertising: 1994-2003", Journal of Advertising, vol. 35, issue 3, pp.143 – 163, 2006.
[6] Adobe Inc., https://www.adobe.com/creativecloud.html
[7] New Research, "Why do people unsubscribe from email newsletters?", https://lab.getapp.com/new-research-getdata-why-do-people-unsubscribe-from-email-newsletters/
[8] D.P. Bertsekas, J.N. Tsitsiklis, Neuro-Dynamic Programming, Athena Scientific, 1996
[9] S.Zhang, R.S. Sutton, "A Deeper Look at Experience Replay", 2017, https://arxiv.org/abs/1712.01275
[10] L. Chittenden and R. Rettie. An evaluation of e-mail marketing and factors affecting response. Journal of Targeting, Measurement and Analysis for Marketing, 11(3):203--217, 2003
[11] R. Gupta, Guanfeng Liang, Hsiao-Ping Tseng, Ravi Kiran Holur Vijay, Xiaoyu Chen, and Romer Rosales, "Email Volume Optimization at LinkedIn", "22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)", pp.97–106, 2016
[12] K. Yang, J.H. Min, and K. Garza-Baker, "Post-stay email marketing implications for the hotel industry: Role of email features, attitude, revisit intention and leisure involvement level", Journal of Vacation Marketing, Vol. 25(4) pp.405–417, 2019
[13] K. Jaidka, T. Goyal, N. Chhaya, "Predicting Email and Article Clickthroughs with Domain-adaptive Language Models", Proceedings of the 10th ACM Conference on web science, pp.177-184, 2018
[14] K. Litinthong, C. Liu, and X. Liu, "An Empirical Study to Understand the Effect of Email Marketing on Consumer's Online Buying Behavior in a Developing Country, "In Proceedings of the 2019 3rd International Conference on Management Engineering, Software Engineering and Service Sciences, pp. 179–183, 2019
[15] Z. Moshe, T. Don, G. Yuval, "Does color in email make a difference?", Communications of the ACM, Vol 49, Issue 4, pp. 94-99, 2006
[16] A. L. Micheaux, "Managing E-Mail Advertising Frequency from the Consumer Perspective," Journal of Advertising, vol. 40, *(4)*, pp. 45-65, 2011.
[17] J. W, Christopher and D. Peter, "Q-learning. Machine learning", vol. 8(3-4), pp. 279–292, 1992
[18] V. Mnih, K. Kavukcuoglu, D. Silver, David, A. Graves, A. Alex, I. Antonoglou, D. Wierstra, M. Riedmiller, Playing Atari with Deep Reinforcement Learning, 2013
[19] H. V. Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning double Q-Learning", in Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16), pp. 2094–2100, 20.