Machine Learning in Marketing Case Study

Why you are more ready for
Machine Learning than you think

Vidora

Table of Contents

A Letter From Vidora's CEO

Thank you for downloading this case study on Marketing and Machine Learning. My name is Alex, and I'm the CEO and Founder of Vidora. I wanted to provide a brief overview of who we are, why we created this case study, and how excited we are about the opportunities for marketers to leverage Machine Learning!

At Vidora, our mission is to make Machine Learning accessible to teams that historically did not have the ability to use Machine Learning. In working with dozens of companies over the years, we found that marketing teams specifically are ready to take advantage of Machine Learning - often more ready than the marketing teams themselves realize they are. Why are they ready? Some reasons include:

- Marketing teams continue to be more data driven
- Data driven teams often capture customer data and use that data to optimize campaigns
- Data driven teams often measure campaign performance

Machine Learning leverages these same foundational pieces of data capture, data usage, and campaign measurement which marketing teams use daily. The difference is that Machine Learning can yield vastly better performance than more traditional data-driven techniques while maintaining existing marketing workflows.

This case study highlights some of the common steps we take in working with businesses to enable Machine Learning. Our ambition in this case study is to convey how easy it is for marketing teams to leverage Machine Learning in their existing workflows and to realize benefits today.

Best,
Alex Holub

This Quarter's Goal: Testing Machine Learning

Let's explore how marketing teams are able to begin driving tangible ROI from Machine Learning in a short amount of time. We'll follow a a typical hypothetical consumer marketing team at the start of a new quarter. When discussing the consumer marketing team goals for the quarter, a theme emerged around the use of data to drive KPIs. Specifically, the team expresses a desire to do more with the data they have in order to improve the KPIs around their marketing campaigns.

Mary, who leads the digital marketing efforts, suggests that Machine Learning will be able to help drive some of the KPIs around campaigns. Mary used Machine Learning at her previous company and has experience with the various ways it can help her team. With consensus amongst the group, a goal for the quarter is set

Team Goal: Test Machine Learning solutions and determine if they improve marketing campaign KPIs

Wanting to ensure the team is focused throughout the quarter, Mary outlines two specific ways the team can explore Machine Learning.

- **Campaign Targeting:** Can Machine Learning help us identify the best audience for each specific campaign?

- **Campaign Actions:** Can Machine Learning help us identify the best action to take to maximize the ROI of a campaign?

Over the next 3 months, Mary will work with her team to test whether Machine Learning will improve their targeting and campaign actions. The results from this analysis will determine if Machine Learning will be a key tool for the team in coming quarters.

Vidora Aside: Why Marketing Data is Ripe for Machine Learning

Why are we focused specifically on marketing teams for this case study? There are a couple key reasons. First, it turns out that typical goals for marketing teams are very well aligned with the results of Machine Learning - namely predicting future behaviors. Second, the available data to marketing can be easily used within a Machine Learning context.

Most marketing is focused on engaging with customers in order to influence their behavior in the future. Without using Machine Learning, teams are often left manually analyzing past behavior in order to make assumptions about what will occur in the future. But, most marketing questions are focused less on what users did in the past and more on what users will do in the future. For example, consider choosing a target audience for a campaign. Choosing this campaign audience is really asking a question about *future* user behavior:

- What audience is most likely to purchase?
- What audience is most likely to churn?
- Etc.

Why is marketing data amenable to Machine Learning? Marketing data is typically captured as analytics events showing customer behaviors over time. This format of data is ideal for answering questions about future user behaviors. In summary, Machine Learning enables marketing teams to use the data they already have, and ask questions they are already asking, while dramatically improving their performance.

Improve Targeting: Predictive Machine Learning

Mary decides to begin her work testing the efficacy of Machine Learning with the team focused on the Churn Prevention campaign that is currently running.

The team's strategy is to target user groups they feel are the most at risk of leaving the service. The team has changed the at-risk target audience definition numerous times, but as it stands today, the campaign is targeted at:

- Users who were last seen >2 weeks ago
- Users with 0 Purchases, or where the last purchase was >1 month ago
- Users who have not been sent this campaign before

Improve Targeting: Predictive Machine Learning

Mary begins by asking two questions which Machine Learning is ideally suited to answer:

- How do we know this is the audience most at risk of churn?
- Is there additional data which can make this target audience more precise?

Mary is not going to answer these questions by analyzing past data and coming up with new rules. Instead, Mary is going to use Machine Learning by asking a predictive question: “Which users are most likely to Churn in the next week?”. She is going to train her Machine Learning prediction using all data the team is capturing on customers, which will allow the Machine Learning model to find the behaviors and attributes that are most predictive of future churn. This provides multiple advantages over the previous at-risk cohort targeting strategy:

Use All Your Data

This model uses all data available to create this prediction. **Meaning**: even data you might not have thought would provide value will be used to improve the accuracy of the prediction.

Control Audience Size

This prediction allows you to rank your customers by their probability to convert. **Meaning**: the size of your audience can be controlled by which range of probabilities you want. (e.g. >50%, >90% likely to convert)

Continuous Learning

New data can continuously be sent to the model to ensure an always up to date prediction. **Meaning**: as user behaviors change over time, the prediction automatically updates accordingly.

With this new prediction, Mary can now update her campaign targeting to utilize all data available. The predictive framework will pick the best data to use in order to generate the most accurate prediction.

Old Targeting

- Users who were last seen >2 weeks ago
- Users with 0 Purchases, or where the last purchase was >1 month ago
- Users who have not been sent this campaign in >1 month

Predictive Targeting

- Customers who are in the top 10% predicted to churn based on the Machine Learning Model using all available attributes

Vidora Aside: How Propensity Models Predict Future Behavior

Above, we briefly talk about Machine Learning predictions that take existing analytics data and use it to predict the likelihood a customer will do an action in the future. But how does this actually work?

1. **Analyze Past Behaviors:** Computers can analyze vastly more data than humans can, meaning all data can be used in targeting, not just a subset

2. **Learn from Past Data:** Predictive Models take the vast amounts of past data and find the behaviors that are most predictive of future actions

3. **Predict Future Actions:** Based on learnings from past data, each user will be assigned a likelihood to complete the action in the future

The benefit of predictions are not limited to he example explored in this case study. Any campaign that has a clearly defined goal can benefit from this process. By using propensity modeling to predict the future behavior of your customers, you can immediately start being more efficient with your data and obtaining answers to questions which leverage all existing data.

Campaign Actions: Matching the Best Action with the Right Audience

Mary is also working with her team on a new email campaign targeted at preventing Churn. This churn redaction campaign is a 5% off email offer designed increase purchases from the Highlighted Item of the Week on the site.

The A/B Test that was recently run revealed that the 5% discount offer increased purchases by 10% over the control group which was not offered the discount. This was a great result for the team, as the discount has been shown to work!

However, this test left many unanswered questions around who exactly the discount should be sent to. The team is asking Mary whether money being left on the table?

Campaign Actions: Matching the Best Action with the Right Audience

Mary realized there are many examples of customer behaviors which could be negatively impacting revenue including:

- Some of the users who purchased with the 5% discount may have purchased anyway, in which case the result would be losing 5% of revenue that would have been generated.
- Some users may only be purchasing because they received the discount, in which case the marketing team is directly generating net new revenue.

So Mary and the team set out to use Machine Learning to see if they can predict the best action to take for each user, specifically who should receive the discount email and who should not receive the discount email. Given the A/B Test was split into two random groups, this means the Machine Learning model can learn from two sets of behaviors. These learnings can then be applied to the rest of the customer base in order to predict the likelihood of purchase in both the scenarios where the user receives a discount email and when the user does not receive an email.

The resulting predictions can help Mary understand the specific effect sending this email will have on each individual customer. With the result being 4 distinct user groups

Customers to Target

Persuadables:
Those customers who were not predicted to purchase without the email discount, but are predicted to purchase with it. This is the exact target audience Mary and the team want for their email campaign.

Customers to Not Target

Lost Causes:
Those customers who are not predicted to purchase in either scenario.

Sure Things:
Those customers who are predicted to purchase even without the email discount.

Disuadabels:
Those customers who are predicted to be less likely to purchase if sent the email.

These 4 groups provide a perfect framework for Mary to decide who will receive the discount email. Customers who will directly use this email to purchase when they might not have without the email will be the future target audience for this campaign. Mary can now optimize the revenue for the company with this marketing campaign.

Vidora Aside: Uplift Modeling

At Vidora, we believe that uplift modeling represents the next big frontier for data-driven marketing teams. By using uplift models, marketers can move beyond predicting each customer's future and toward finding actionable strategies for positively influencing them.

What is uplift modeling?

Uplift modeling is a prescriptive Machine Learning technique which predicts how each customer is likely to respond to a marketing action. Based on these predictions, marketing teams can determine which customers should be targeted in order to maximize the ROI of their campaign.

Defining an Uplift prediction involves specifying two inputs:

- **Intervention**: Which marketing action(s) would you like to measure the impact of?
- **Outcome**: What is the conversion event that your intervention is meant to influence?

For example, if we are considering offering a discount (the intervention) with the goal of driving customer transactions (the outcome), the Uplift prediction for Customer A might tell us: "A discount will increase Customer A's probability of purchasing by 10% compared to if we did not offer the discount."

In order to create an Uplift model, we must have data from a randomized control trial in which a random set of customers were targeted with the intervention (the "treatment group") and another random set of customers were not targeted (the "control group"). An Uplift model uses this information to differentiate between four types of users (see previous page). If the model is able to accurately identify Persuadables and Disuadabels (sometimes referred to as "Sleeping Dogs"), marketers can determine which customers to target in order to drive incremental conversions with a marketing campaign.

Testing Machine Learning: Real World Testing

This is the first time many of the members of Mary's team have directly worked with Machine Learning. For this reason, ensuring everyone is confident in the predictions is a key step in adopting Machine Learning for the long term.

To help with this, Mary is measuring and reporting on success metrics at three separate points in the Machine Learning Process:

1. **Model Training:** How well did the model do at predicting data it already has access to?

2. **Prediction Tracking:** How accurate were the predictions when compared with real world conversions that occurred after the predictions were made?

3. **Real World Effectiveness:** When predictions show positive results in #1,2, the team then starts to use the predictions in actual campaigns and measure the outcomes

What she and the team found was promising for both (1) and (2). Model Training and Prediction Tracking showed high accuracy. This was a good sign because it meant the team had enough data from which the model could best learn the predictive behaviors.

Next Mary turned to analyzing the Real World Effectiveness of the campaigns on the Weekly Highlighted Item of the Week discount email. Mary saw that targeting only the persuadable group resulted in much higher overall revenue for the company.

As for the Churn Prevention campaign, the team has decided to continue to A/B Test the best use of this prediction. The predictions were shown to accurately predict those users who will churn, but the team is still finding the best threshold point at which a user is deemed at risk enough to warrant the outreach.

Overall the results were great. And Mary is working with the team to integrate Machine Learning into an increasing number of campaigns.

Vidora Aside: Are you ready for Machine Learning?

Vidora often works with marketing teams taking their first path on the journey of using Machine Learning. The use cases outlined above are intended to show how existing campaigns can be improved with Machine Learning using data that is already captured by the team.

We hope to convey a few key takeaways:

- As a marketing team, you probably already have access to the data necessary for machine learning
- With just a few months of data, you can create Machine Learning predictions about your customers
- If using the right platform, getting started and testing Machine Learning can be quick and easy

We encourage all marketers to take a look at the data you are currently capturing, and the campaigns you are currently running, and to think about how Machine Learning can help more effectively put your data to better use.