Piloting A Predictive Analytics Vendor

An Enlyft White Paper

Enlyft delivers actionable sales and marketing intelligence to help companies understand who their most promising customers are, why they are the right customers, and what the most effective messaging is to approach them with.

enlyft.com

enlyft

Piloting A Predictive Analytics Vendor

The process of selecting a predictive analytics vendor can be challenging. Far too often predictive analytics vendors cherry-pick results, highlighting only their most impressive ones.

Participating in a pilot with a predictive analytics vendor allows you to see the results of the vendor's services immediately and at low cost. You'll not only learn what it's like to work with the vendor, you'll also be able to assess its overall performance compared to your existing solutions, which may be nonexistent, or may include firmographic targeting, technographic targeting, or other predictive modeling capabilities. This guide is intended to assist you in successfully carrying out a pilot project with a predictive analytics vendor.

You have to be able to establish a controlled process that is reliable.

I think Enlyft can help there because they have plenty of experience of doing rigorous A/B testing.

John Jenkins, Microsoft

Establish high-level guidelines for the pilot

Even before the start of the pilot, you should work with the vendor to decide on high-level guidelines.

First, you should engage in a high-level discussion on objectives for the pilot, as well as evaluative success criteria. Next, you should agree on the duration of the pilot, the number of accounts you wish to allocate to your sales team, and the number of performance checkpoints you'll have (you should have ongoing checkpoints throughout the pilot).
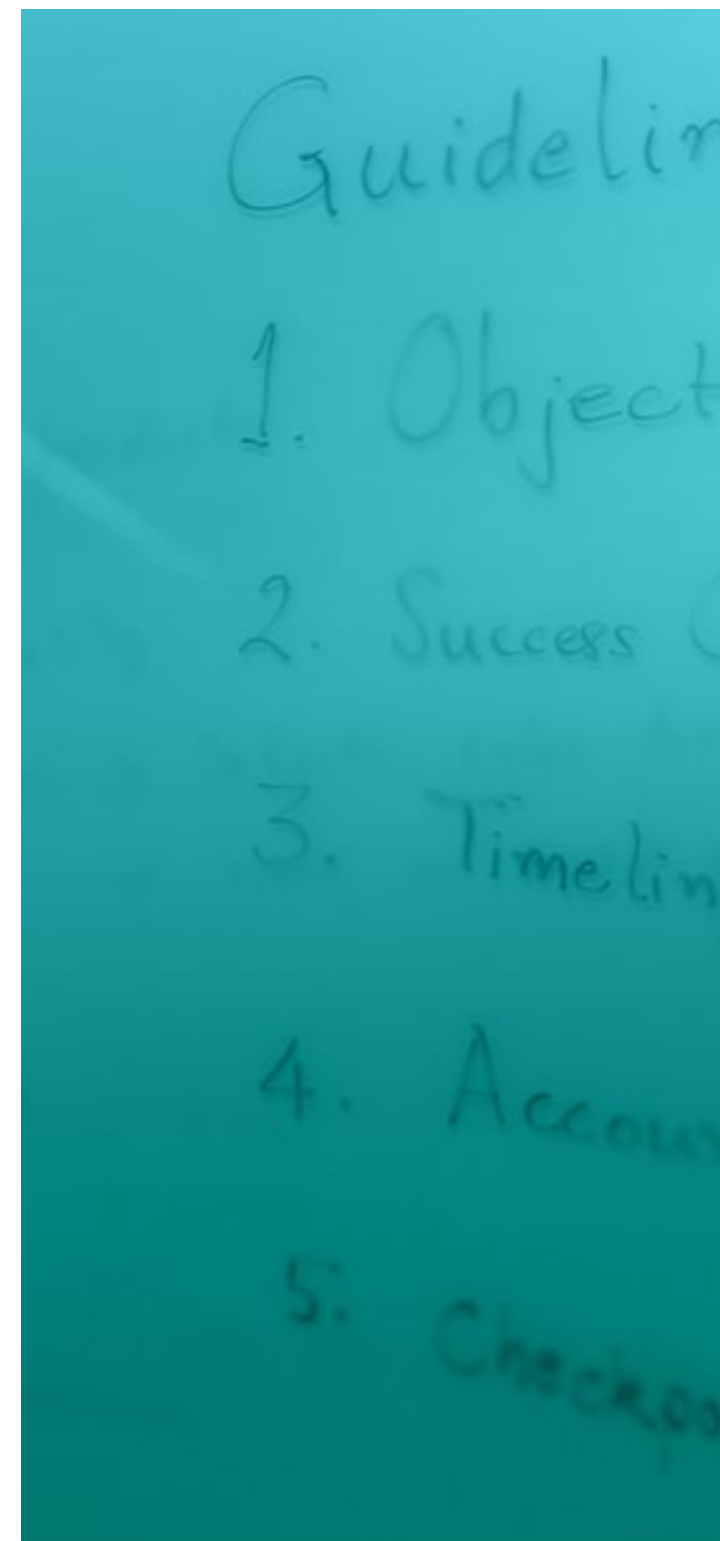
Identify the primary objective and key success metrics

Your vendor should work with you on selecting an objective for the pilot. Ideally, you should select an objective that is challenging, applies to many potential accounts, and for which there is considerable data available, both in the public domain and in your internal sales and marketing systems.

It's important to establish your key success metrics at the onset of a pilot. While your ultimate metric may be increased sales revenue, it's important to identify multiple metrics that align with different stages of the sales and marketing funnels. These metrics may include Marketing Qualified Lead (MQL) volume, MQL to Sales Qualified Lead (SQL) conversion rate, MQL to Opportunity conversion rate, average deal size, average win rate, and Opportunity to Win conversion rate. These intermediate metrics allow you to get a feel for the model's benefit without needing to wait for the entire sales cycle to complete.

Ideally, you should be able to see statistically significant results with the metric within 4-6 weeks. Enlyft has found that MQL to SQL conversion is a metric that works well as the primary intermediate milestone. It often results in statistically significant results in just a few weeks and standard conversion rates from SQL to Closed/Won are often well know.

When evaluating success metrics, it's critical to consider sales cycle variability. If you have considerable variability in your cycles results (from month to month or quarter to quarter, for example), we recommend comparing the vendor's recommendations with your existing methods simultaneously. This will allow you to carry out an A/B test and prevent the vendor's results from appearing artificially inflated or deflated. This will be discussed below.
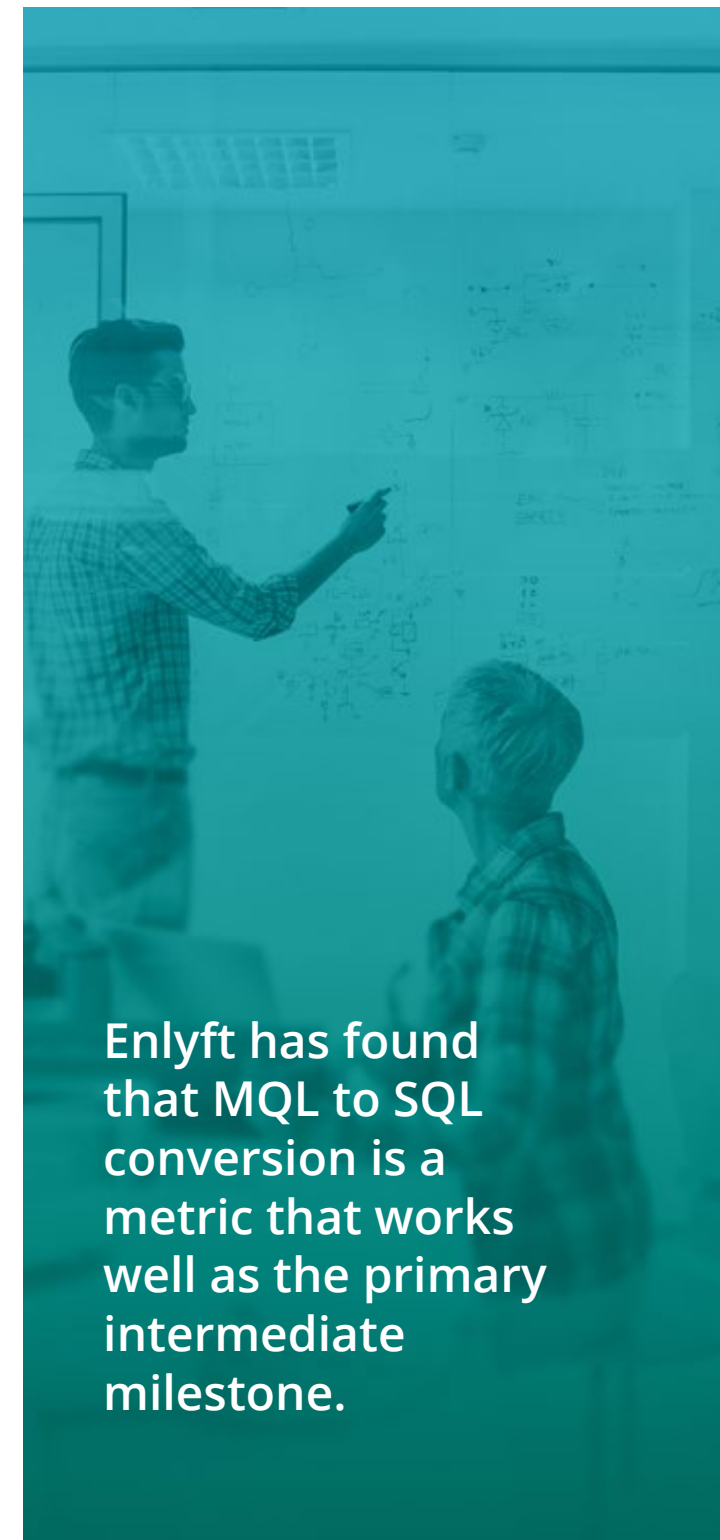
Enlyft has found that MQL to SQL conversion is a metric that works well as the primary intermediate milestone.

Select win/loss accounts to train the model

It's important to carefully select the win and loss accounts that will be used as part of training data. As the name implies, training data will be used to "train" the model, enabling it to find accounts that are more similar to your wins and less similar to your losses.

While machine learning models can work effectively with minimal data, the more data provided, the better. Enlyft recommends providing training data of at least 500 won accounts and 500 lost accounts. This volume of training data enables us to withhold 20% of the dataset to use as "holdout" set. This testing data will be used to estimate the model's performance and guide adjustments to the model.

It's also important to select account loss data carefully. Ideally, you should exclude any outliers from the training set. These outliers should be otherwise desirable accounts that had a high win likelihood but were lost unexpectedly, perhaps due to an anomaly or abnormality. It's helpful to discuss particular details and nuances associated with account losses with the vendor. For example, if you are looking to gain a competitive advantage and win more accounts when going head-to-head with specific competitors, discuss this with your vendor. In this case, you'll likely want to include data pertaining to desired accounts that were lost in head-to-head competition in the training set. The model will be able to uncover the subtle reasons why specific accounts were lost

Enlyft recommends providing training data of at least 500 won accounts and 500 lost accounts.
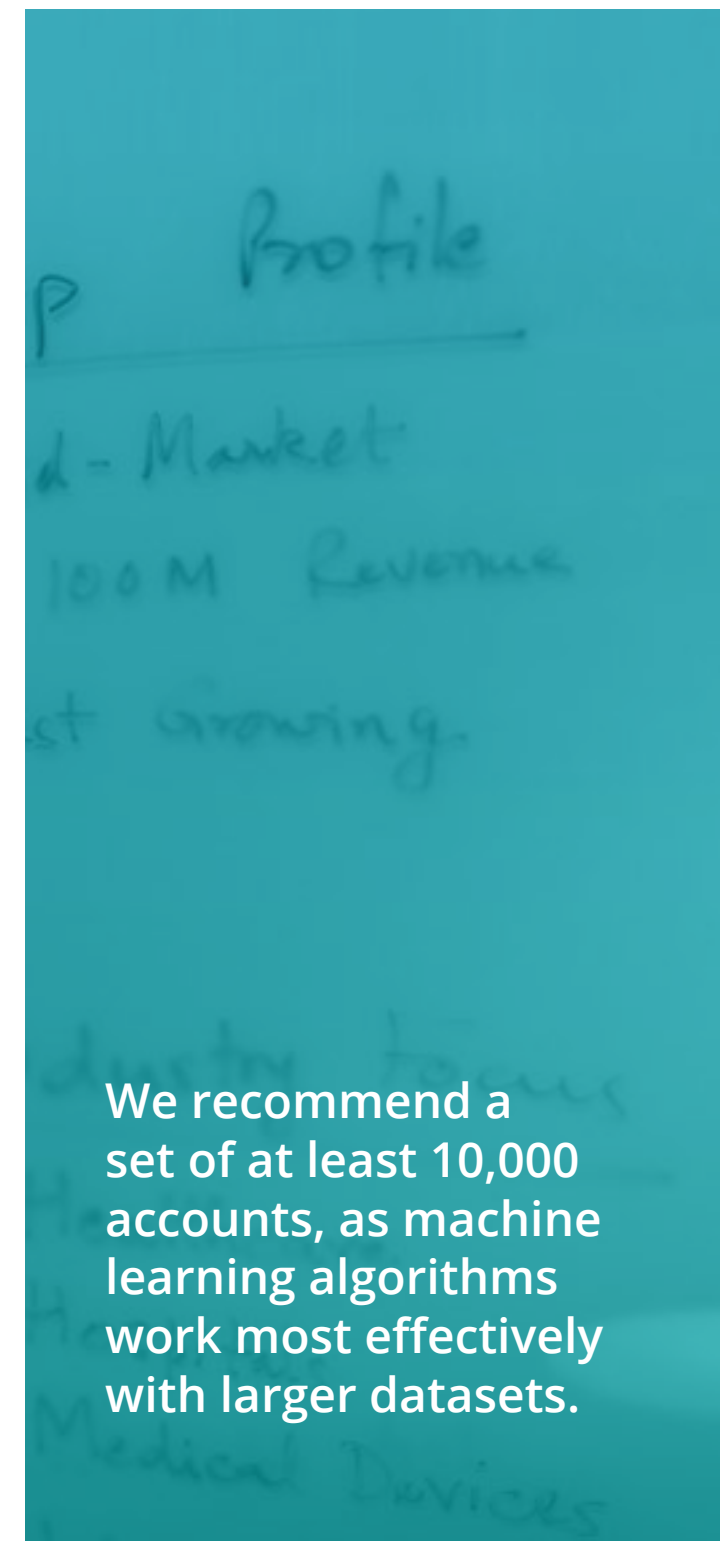
Select a set of accounts to prioritize

After you've selected the training data, you'll need to determine a set of accounts to be prioritized by the model. We recommend a set of at least of 10,000 accounts, as machine learning algorithms work most effectively with larger datasets.

When selecting a set of accounts, you should zero in on a particular group of accounts that share similar characteristics. For example, if your total addressable market comprises of medium and large-sized enterprises based in the United States, you may want to include only medium-sized enterprises, only largesized enterprises, or only medium and large-sized enterprises located in the Midwestern United States.

Narrowing your focus will enable you to obtain a true measure of your success and obtain a good feel for how well the model will perform when you scale up to target your entire market. If you don't have an existing set of accounts that you'd like to use, some vendors are able to generate a set of accounts from those contained in their database (based on desirable firmographic and technographic criteria).

We recommend a set of at least 10,000 accounts, as machine learning algorithms work most effectively with larger datasets.

Determine which internal sales and marketing data to use

Your internal sales and marketing data can play a significant role in determining best fit accounts.

Typically, most of this data will consist of account specific data that is housed within your CRM system. Consider all the types of account data you routinely collect (e.g., purchase history, sales interactions, social media engagement, past white paper downloads, and webinar attendance history). If there is a specific ideal customer feature set that you would like to target, discuss this with the vendor (even if you don't have any pre-existing data pertaining to specific ideal accounts in your CRM system).
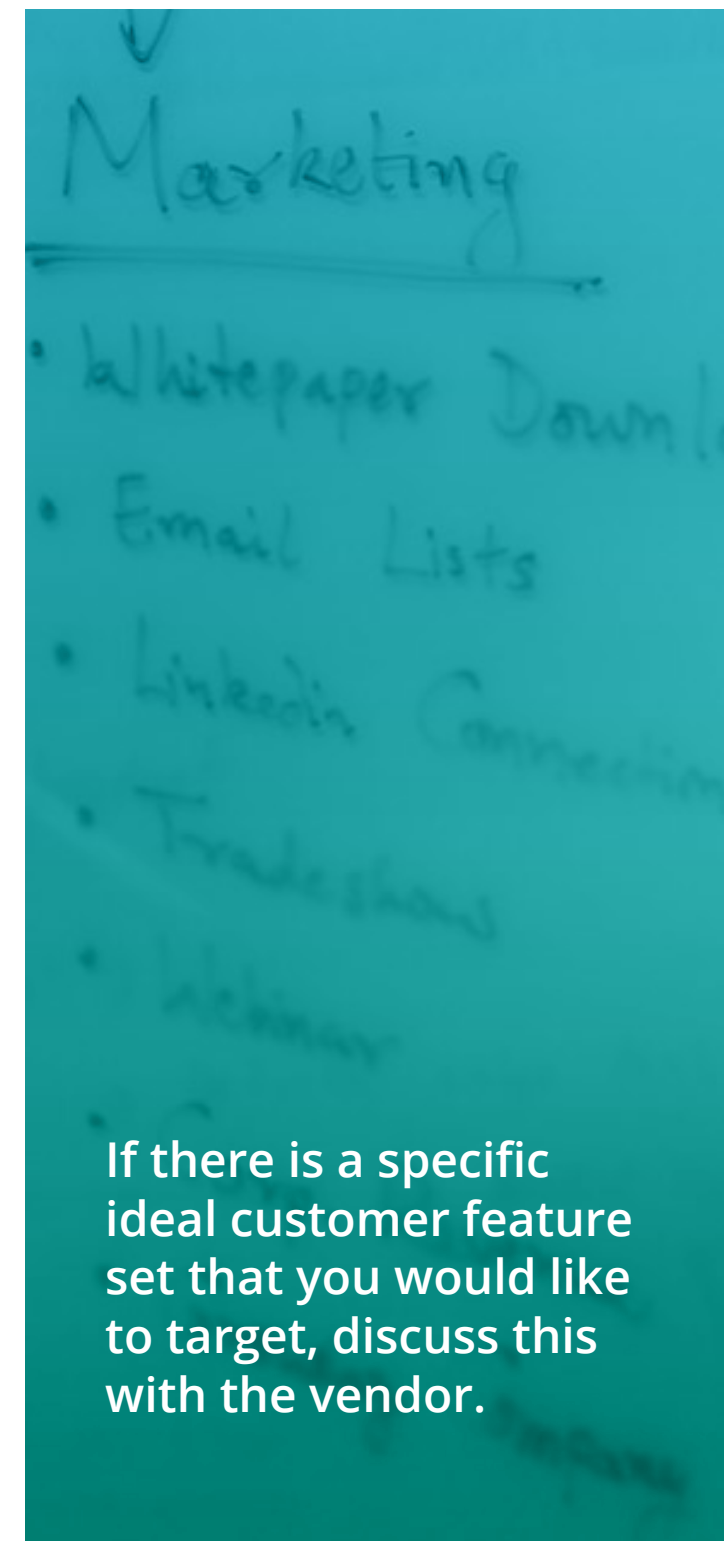
If there is a specific ideal customer feature set that you would like to target, discuss this with the vendor.

Determine how many accounts to allocate to your sales team

In order to decide how many accounts to allocate to your sales team, you'll need to determine the number of accounts that your sales team can reasonably process in a given time frame. Consider, for example, that your sales team is able to move approximately 800 accounts to the Sales Qualifying Lead (SQL)stage in a reasonable time frame. You should then use about 400 account recommendations from the vendor (comprised of approximately 300 high propensity accounts and approximately 100 low propensity accounts) and another 400 accounts from your existing method of selecting accounts (once again comprised of approximately 300 high propensity accounts and approximately 100 low propensity accounts).

In compiling the final list of the 600 high propensity accounts (300 from the vendor and 300 from your existing selection method), you'll need to determine the number of unique accounts contained in each list. If the list of high propensity accounts generated by each method includes more than 300 unique accounts, simply select the top 300 from the list. If, however, the list has fewer than 300 unique accounts, you'll need to supplement the unique accounts with additional accounts so that the final list from each method comprises 300 accounts. You can accomplish this by developing a list of accounts generated by both methods and then selecting from this list.

Selecting the low propensity accounts is simple. Look at the bottom of the lists of accounts generated by each method and select the lowest propensity accounts. If there are fewer than 100 unique accounts, adhere to a similar methodology as described above to obtain 100 total accounts. In the case of losses, there typically isn't much overlap between the methods.

Note that the number of low propensity accounts is significantly less than the number of high propensity accounts. Given that low propensity accounts tend not to perform well, you should select as few accounts as possible, while ensuring you select a sufficient number to see a statistical difference between the performance of the high and low propensity recommendations. Once this process is complete, you'll obtain 800 accounts to pass off to your sales team.

COMPILE ACCOUNTS FOR YOUR PILOT
Create your accounts list from an even split of vendor supplied and your own accounts.

A | Define the number of accounts required.
800 (in our example)

B | SOURCE HALF FROM THE VENDOR
Use 3/4 high propensity and 1/3 low propensity accounts.

300 + 100 = 400
High   Low   Total

C | SOURCE THE OTHER HALF USING YOUR NORMAL APPROACH
Use 3/4 high propensity and 1/3 low propensity accounts.

300 + 100 = 400
High   Low   Total

D | COMBINE THE LISTS FROM BOTH SOURCES
Creating an accounts list to pilot

400 + 400 = 800
Vendor   Your own method   Total

NOTE: Green numbers are our example figures

Understand the model and recommendations

In addition to providing a list of prioritized accounts, the vendor should be able to provide you with an explanation for the model generated, as well as recommendations.

You should be able to understand how the model works, including the key factors that represent your Ideal Customer Profile (these factors should be based on your internal sales and marketing data as well as external data provided by the vendor). Ideally, the vendor will supply information pertaining to why each account has a high or low propensity. Understanding the "why" is critical, not only in terms of effectively communicating the value of the model internally, but also in terms of empowering sales team members to engage in more meaningful interactions with target accounts.
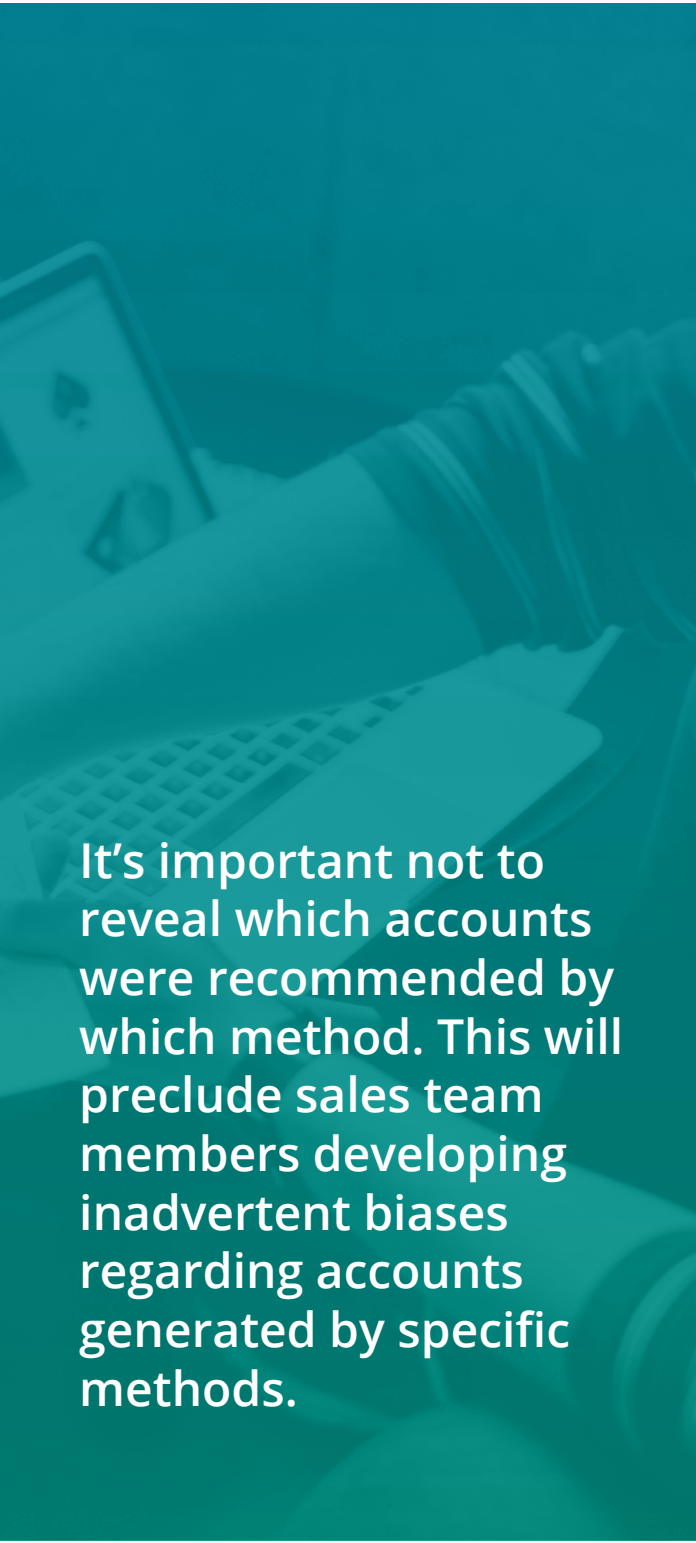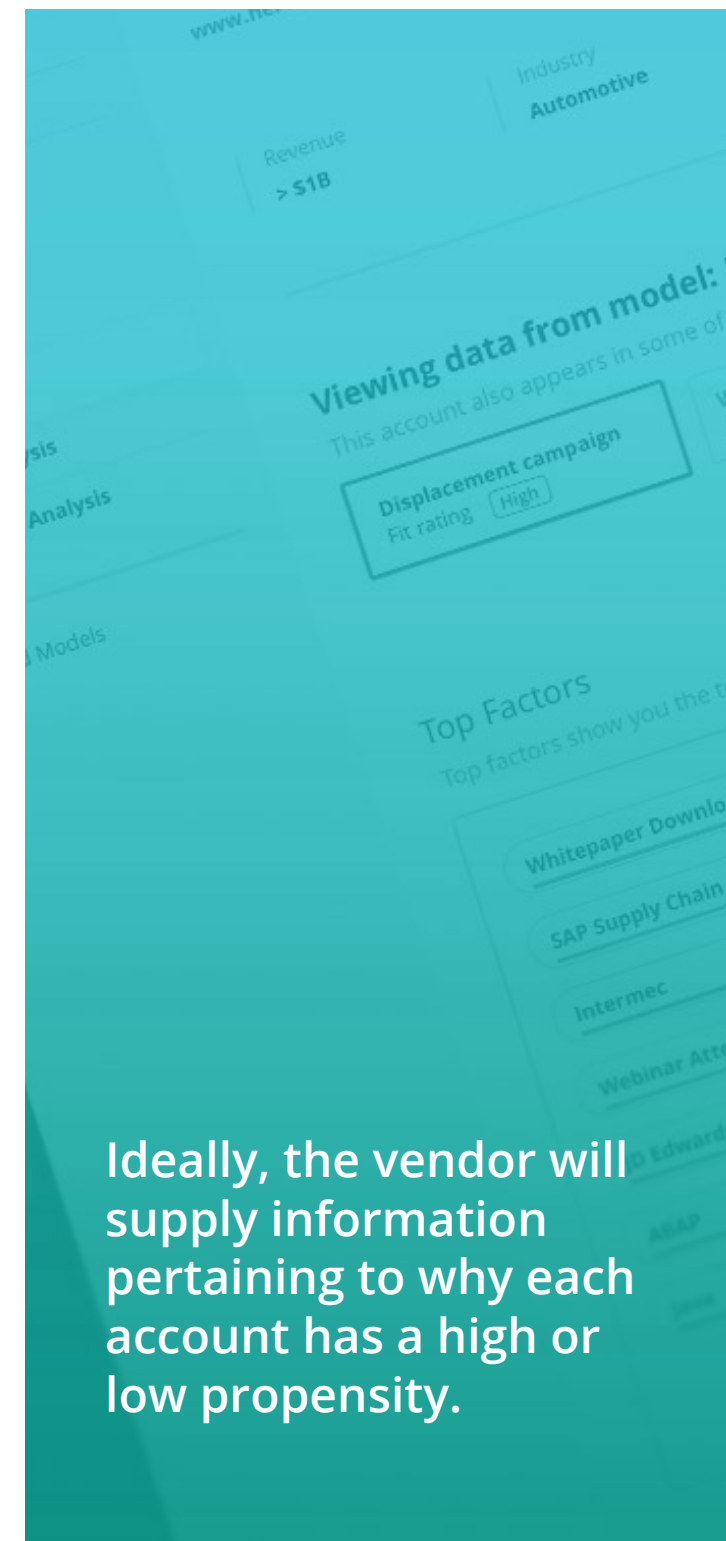
Ideally, the vendor will supply information pertaining to why each account has a high or low propensity.

Run the pilot in the field

You now need to hand off the prioritized account lists to your sales team. In the case that you plan to pass off 800 accounts to your sales team, you should combine the 400 accounts recommended by the vendor with the 400 accounts generated by your existing methods and then randomize the ordering of the combined list.

After the account lists are combined, you'll want to randomize the list and pass it off to your sales team in a way that minimizes the likelihood of sales team members giving unfair or biased treatment to certain accounts. It's important not to reveal which accounts were recommended by which method. This will preclude sales team members from developing inadvertent biases regarding accounts generated by specific methods. Additionally, if a vendor assigns a specific rating or number to each account, avoid disclosing this information to your sales team. This will prevent sales team members from investing more effort to target higher-rated accounts. All accounts should be treated equally.

It's important not to reveal which accounts were recommended by which method. This will preclude sales team members developing inadvertent biases regarding accounts generated by specific methods.

Evaluate the results at every checkpoint

You should have multiple checkpoints throughout the pilot. At each checkpoint, you'll be in a position to evaluate results according to your chosen metrics.

In order to determine each method's performance, you'll need to perform a series of simple calculations. First you'll need to determine four success rates: one for the high propensity accounts generated by the vendor (300 accounts in the example above), one for the high propensity accounts generated by your existing selection method (300 accounts in the example above), one for the low propensity accounts generated by the vendor (100 accounts in the example above), and one for the low propensity accounts generated by your existing selection method (100 accounts in the example above). In order to calculate each conversion rate, simply divide the number accounts that resulted in wins according to your predefined success metrics by the total number of total accounts in each group that was passed off to your sales team. The win rate is a useful metric for comparing the relative performance of each method.

After you've completed your calculations, you should generate two ratios. These ratios will give you a good indication of each method's ability to separate the valuable accounts from the invaluable. You can calculate the ratio by dividing the success rate of the high propensity recommendations by the success rate of the low propensity recommendations. You should perform this calculation for the accounts generated by the vendor, as well as the wins generated by your existing method. The better performing method will exhibit a higher ratio and provide better results as per your pre-determined success metrics.

CALCULATE THE SUCCESS RATE

For each group of accounts. Divide the number of accounts that reached the SQL stage.

SQL   Resulting success rate

$$\frac{30}{300} = 0.1$$

Total No of accounts in group

Repeat the same approach for each of the four sets of accounts. This gives you an individual win rate for each set of accounts

CALCULATE THE SUCCESS RATE

For each source of accounts. Divide the success rate of the high propensity recommendations by the success rate of the low propensity recommendations.

Success rate of high propensity accounts   Resulting win rate ratio

$$\frac{0.1}{0.02} = 5$$

Success rate of low propensity accounts

Repeat the same process for each source of accounts to give you a comparable figure for both.

NOTE: Green numbers are our example figures

Example Pilot Timeline

The timeline for a pilot can vary significantly based on your sales cycle, as well as your ability to generate win/loss and other data required by the model.

Ideally, the time required by the vendor to perform key activities, including data cleaning, developing the model, and analyzing the results, should be fairly predictable and quick. Outlined below is the timeline associated with a recent pilot Enlyft performed for a medium-sized B2B SaaS company (after the goal of the pilot was established).

1 Discuss metrics and details of sales processes

During an initial 1.5 hour meeting with the client, Enlyft confirmed the use cases for the model and agreed on the key success metrics that would be tracked throughout the pilot. The client's sales processes were discussed and the relevant sales and marketing data that would be used to generate the model were agreed on.

2 Collect data

Over the course of two weeks, the MQL to SQL conversion rate, target account, and sales and marketing data were collected and provided to Enlyft. Enlyft's data scientists cleaned and reviewed the data. The A/B test plan was established and two lists of accounts were provided to the client. The first file included a list of accounts and the associated rankings (e.g., high versus low propensity). The second file included a randomized list of the recommendations without the associated propensity rankings. The client distributed this list to its sales team to work through.
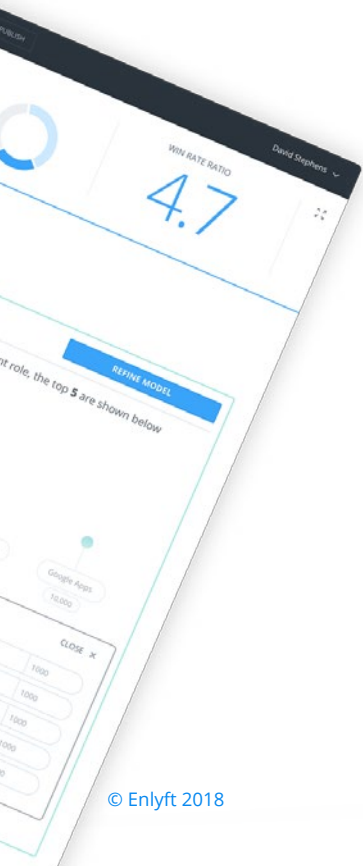
3 Review results at first checkpoint

After about three weeks, sufficient progress was made in terms of the early funnel metrics that the results could be analyzed. Within a couple of days, Enlyft generated a report of results (concluding that the high propensity recommendations significantly outperformed the low propensity recommendations).

Learn more about Enlyft

Visit us at enlyft.com or email us at sales@enlyft.com

enlyft