A Work Project, presented as part of the requirements for the Award of a Master's degree in Business Analytics from the Nova School of Business and Economics

A ML-based customer reference recommendation system to optimize the sales process at a B2B software company

Theresa Isabel Kruse

Work project carried out under the supervision of:

Qiwei Han

17-12-2021

Table of Contents

<u>Figures</u>

<u>Tables</u>

Appendices

Disclaimer: Due to data privacy issues with the addressed company, this published paper is censoring some of the information given.

Title: A ML-based customer reference recommendation system to optimize the sales process at a B2B software company

I. Abstract

As account-based marketing and customer engagement are major sales drivers in B2B companies, this paper aims to examine and improve the customer referencing at a B2B software company. Customer referencing allows prospective customers to interact with existing customers and learn about their experience with the company's software. Currently, prospects and possible advocates are matched manually, resulting in a time- and resource-intensive procedure. The goal of this paper is to propose a machine learning-driven recommendation system to match prospect requests with the company's existing advocates based on their similarity.

Keywords: Business Analytics, Data Analytics, Machine Learning, Multilabel Classification, KNN, Account-based Marketing, Customer Engagement, Sales Reference, Customer Reference, Process Optimization

1. Introduction

This paper identifies a solution to recommend best matching advocates for customer reference requests made by prospect customers at a European software company. Before, when a reference request is made by a prospect company, matches are manually searched for, which slows down the matching process. To solve this problem, best matches are generated using a machine learning (ML)-driven approach and the result is pushed into an application, called Reference Dashboard. Two lists are proposed for matching each request: One shows best matches among customers that previously showed a willingness to engage in advocacy and the other one includes best matches among all possible customers.

First, a theoretical overview of customer reference systems is given, using a funnel approach by introducing account-based marketing followed by customer engagement. This enables the reader to understand the importance of a customized marketing and sales approach in business-to-business (B2B) companies. In a next step, this is narrowed down to explaining why customer references provide highly transparent and trustworthy validation of a product or service. Further, the paper describes the technical approach taken to identify best matches among existing customer accounts for reference requests. In a first step, use cases that the prospect is likely to implement using the software were predicted and then used as additional features to match the prospect with existing accounts based on similarity.

2. Sales Customer Reference Process

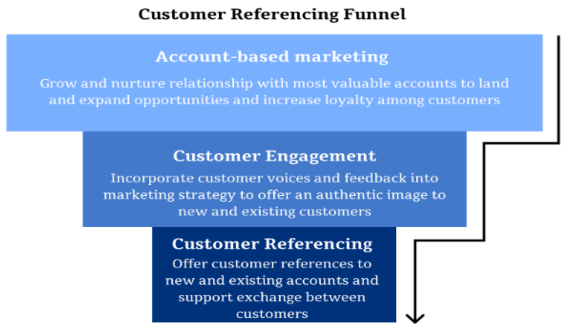

Figure 1: Sales Customer Reference Funnel

Aiming to explain what a good reference is and how the software company is generating advocates out of customer accounts, a funnel approach, shown in figure 1, is used. Starting with account-based marketing as a customized way of targeting key customers and turning them into advocates, the value of customer engagement and its benefits to B2B firms, is discussed. Finally, the customer reference process, the focus of this paper, is presented as a form of customer engagement.

3.1. Account-based Marketing

The number of approachable companies for B2B businesses can be limited (Stadelmann, Pufahl, and Laux 2020), which forces these businesses to make efforts to grow existing customers' loyalty and convince them of their superiority to their competition. Within a potential prospect firm, "the number of individuals needed to convince before making a purchase are numerous" (Gene Day and Wei Shi 2020). With modern technology, customers and prospects are far more well informed than some years ago and can easily compare a company to its competitors (Povolna 2017), which causes higher expectations and more complex requests (Andersen et al. 2017). This induces the need for a marketing strategy that customizes the sales and product experience for customers and prospects, which is called Account-based Marketing (ABM).

ABM is a strategic approach (Bacon 2021) to target the most important customers or prospects to land or expand opportunities in B2B companies. "ABM has specific objectives and a narrow audience" (Burgess 2016) and is often referred to as "Flipping the Funnel" when compared to traditional Demand Generation Methods (Vajre 2016) (see Appendix A). This comparison originates from the strategy to first approach a lead that identifies as high potential account and building a relationship with that account from there: “Account-based marketing is like learning a new language. You’re no longer targeting demographics or personas, but rather specific organizations” (Pavan Kumar and Rajasekhar 2020). Marketing, sales, and account management departments must work closely together (Burgess 2016; Bacon 2021) to “focus […] on building relationships with prospects, gaining a deep understanding of their needs, and offering ongoing consultations and solutions” (Pavan Kumar and Rajasekhar 2020). ABM is understood and addressed differently from company to company - “there is no one size fits all” (Bacon 2021).

The *pareto principle*, also referred to as the *80/20 Rule* helps understand why ABM is such an essential approach for B2B firms. It sheds light on the fact that “80% of results will come from just 20% of the action” (Kruse 2016) or in other terms that “20% of customers account for 80% of total profits” (Kruse 2016) of a company. Looking at the total revenue of US$ 677 of the three leading tech companies, Apple, Samsung Electronics, and Alphabet (Ponciano 2021), it is shown that it makes up about 3% of the US gross domestic product (GDP) and 15% of German GDP (The World Bank Group 2021).This proves that focusing sales and marketing efforts on key customers can generate more revenue for the company than penetrating many small ones.

ABM can be implemented on existing accounts to discover new or expand on existing opportunities or on prospective customers to raise their interest and convince them of a product’s superiority over competitors. According to Lucie Povolna, “existing customers possess much larger value than customers which are new to the company and the extension of cooperation with them is essential for maximizing the value of the company” (2017).

Three types of ABM exist: One-to-One, One-to-Few, and One-to-Many (see Appendix B). One-to-One or Strategic ABM focuses on chosen accounts that have a high recurring value for the company. This form of ABM is very resource-intensive, as high investments into research must be made over a long period (Bacon 2020). ABM Lite focuses on similarity-clustered groups of high-value accounts. This is less resource-intensive than Strategic ABM, as resources can be shared among a cluster. The least expensive form is Programmatic ABM, which focuses on a larger number of accounts and is closest to traditional demand generation channels (Bacon 2021).

3.1.1. Account-based Marketing at a B2B software company

At the software company we are looking at, ABM is managed by the field marketing department that is usually for conventional demand generation. According to the Director of Field Marketing (see Appendix C), the company is using a Strategic ABM approach, with responsible very few targeted accounts. Even though this approach is considered more resource- and investment-intensive, the company identified it as a less laborious approach that can be tackled within the existing structures of this department. As stated by the Senior Field Marketing Manager (See Appendix D), each field marketer manages a different region and is "alongside" handling ABM campaigns for key accounts within this area. Thus, ABM has no dedicated role within the company and, according to the Director of Field Marketing (2021) (see Appendix C), a lot more education and training is needed in this area to align understanding and tactics among the team.

Currently, the company identifies ABM accounts by following strict criteria, where either the predicted actual cash value (ACV) or the annual recurring revenue (ARR) of an a customer account must be above a certain threshold (Field Marketing Department 2021). Most accounts approached are existing customers, however, the Senior Field Marketing Manager (2021) (see Appendix D) mentioned that ABM could work very well for potential accounts if the focus was shifted.

Strong coordination between sales, customer success, marketing, and business development teams is needed. According to the Director of Field Marketing (2021), communication among the departments varies from account to account. Thus, a basis must be built that aligns communication, bandwidth, priorities, and resources among the teams (see Appendix C). Another challenge in the expansion of ABM at the company is the goal setting for the field marketing department. The incentive lies in fulfilling a certain pipeline within one year, but ABM must be approached and measured over a much longer timeframe (Senior Field Marketing Manager, 2021) (see Appendix D).

As mentioned above, ABM can be leveraged by companies to build strong relationships with key accounts or corresponding personas and thereby generate more sales at a larger scale. Additionally, this approach improves the buying experience of the customer and strengthens the loyalty and trust towards the company. These motivational drivers can turn customers into advocates who engage in the form of co-marketing and help the company acquire potential customers or advance opportunities. Aarikka-Stenroos and Jaakkola describe this phenomenon of customer engagement: "If customers experienced sufficient value from the use of the solution in its application contexts, they began to act as co-marketers" (2012) The following section throws light on the importance and application of customer engagement and the power of word-of-mouth (WOM).

3.2.Customer Engagement

Customer Engagement is a term used in the context of B2B marketing and sales and is defined by the leveraging of customer relationships in marketing and sales (Harmeling, Moffett, and Carlson 2016; Elina Jaakkola and Alexander 2014). These voluntary customer contributions within companies "occur in interactions between the focal object and/or actors and result from motivational drivers" (Jaakkola and Alexander 2014).

For B2B technology companies, each customer receives individual solutions adapted to their specific needs, which makes it more difficult for these companies to deliver proof of value throughout the sales process. Thus, a positive image and customer satisfaction are vital to demonstrating the value of the product to prospects: “Customer commitment to the business is an important signal to promote joint involvement in developing the business” (Ruokolainen 2007).

While customer engagement is not reduced to the most-valuable customers, ABM drives their engagement to attract more of their kind. Companies can have a large impact on how engagement is leveraged to approach and persuade potential customers and discover new opportunities.

Malthouse et al. mention two types of customer engagement. First, “lower engagement, in which consumers only passively consume content or use very basic forms of feedback”, such as ‘liking’ a post on social media (Malthouse et al. 2013). In higher engagement, the customer actively participates in co-marketing activities, such as testimonials, referrals, or references (Malthouse et al. 2013; Elina Jaakkola, Aarikka-Stenroos, and Kimmel 2014; Bijmolt et al. 2010). Today, online sources are strongly used for the exchange of information about products or services. However, these sources are easily forgeable, which causes offline forms, such as WOM, to be considered more reliable. “The power of WOM is based on the accessibility, credibility, and trustworthiness of the information source” (Jaakkola, Aarikka-Stenroos, and Kimmel 2014).

3.3. Sales Customer Referencing

Sales customer references, which are the focus of this paper, are one form of WOM. Customer referencing occurs when a company reveals its customers to demonstrate the performance and credibility of a product or service (Jing et al. 2020; Zhang 2018; Ruokolainen and Igel 2004). In the following, these customers are referred to as *Advocates,* while the customers requesting this service are called *Prospects*. Prospects can be potential customers or existing customers that want to identify and discover new opportunities within our respective company.

Customer referencing is based on the theory of signaling, which describes the information asymmetry between different parties (Connelly et al. 2011). The theory assumes that "prospective buyers face information asymmetry because they cannot experience the benefits of the product or service until after adoption" (E. Jaakkola and Aarikka-Stenroos 2019) (see Appendix E). Thus, this asymmetry is balanced by letting an advocate signal his or her experience with a product or service to a prospect. Moreover, credibility and the reduction of perceived risk are further benefits of the exchange (Morgado, Mota, and Castro 2015; Salminen and Möller 2006; E. Jaakkola and Aarikka-Stenroos 2019; Blazevic et al. 2013). The general perception is that receiving information from an experienced customer is more authentic than from the company itself: "While evidence of competence presented in a firm's advertising may be viewed as biased, evidence provided via customer reference is seen as more independent because it is from third party" (Zhang 2018).

4. Problem Formulation

To understand the problem and solution that this paper is addressing, it is essential to understand the current sales customer reference process at the respective company, which is outlined in the following section.

4.1. Sales Customer Referencing at a B2B software company

At the addressed B2B software company, sales customer referencing is available to all potential customers or customers who want to discover new opportunities. The main reasons identified for requesting a reference at this company are to learn about functionalities of the software or best practices regarding the implementation (Internal Database, n.d.). Furthermore, references are used to double-check the investment and to convince top management of the purchase (Internal Database, n.d.).

Figure 2: Customer Sales Reference Process at the B2B software company (censored due to data privacy issues)

As shown in figure 2, five parties are involved in the sales referencing process: the prospect, the *Sales Representative* (SR), the *Customer Reference Manager* (CRM), the advocate, and the *Customer Success Manager* (CSM) of that advocate. The prospect is the company that requests the reference, who wants to gain more knowledge about the company and software from an existing customer. The SR's goal is to convince the prospect of investing in the software or to expand on its existing subscription. The advocate belongs to an existing customer account who aims to share knowledge about his or her experience with thproe product. The motivation to share information is mainly driven by enthusiasm about the product, a strong relationship with the assigned CSM, or the willingness to network with similar organizations. The CSM takes care of the advocate account and will reject a reference request if it is not in the best interest of the advocate account. The customer reference process is managed by the CRM, who is part of the *Customer Marketing Department* (CMD) and aims to offer the best reference meeting experiences to both prospects and advocates.

When a reference is requested, the SR contacts the CRM and the information is added in the sales software. The CRM then searches for a fitting request in the *Advocate Finder* (AF), a list in the sales software that stores all advocate contacts and can be filtered by industry, region, and use case. When a fitting advocate is found, the CSM of that advocate is contacted, who can reject or approve the request in consultation with the advocate. If the advocacy request is rejected, the CRM must search for a new advocate until an available advocate is found and the request accepted by the corresponding CSM. The CRM then schedules a meeting between the prospect and the advocate. After the meeting has taken place, feedback on the process is requested from both parties and the CSMs and SRs participation in the reference program is rewarded.

A Customer Marketing Hub (CMH) is to be introduced where all customer marketing activities can be tracked and managed on one platform. This way, customer marketing processes should be made more generic and easier accessible for all parties. The customer reference process is also to be part of the CMH with an optimized process that this paper is outlining. The challenges and problems with the current process and a suggested solution are discussed in the following section.

4.2. The Problem

As can be seen above, the reference process is complex and can contain many loops in cases in which advocacy requests are rejected. Rejections can have many reasons: For example, accounts can be temporarily disqualified from advocacy by being in a renewal phase of their subscription or having engaged in too many advocacy activities in past months. With the current process, these criteria are not easy to identify by the CRM who must frequently reach out to CSMs to identify the company's availability for advocacy. Furthermore, rejections can be due to the CSM or SR deciding that the match is not ideal and not wanting to engage their accounts in any unnecessary activities. This is bound to a lot of research to find good matches for each request. Having only three filters to search for advocates in the AF often throws the same advocates at top rows while some stay hidden at first glance. Because advocates can only be requested once every three months and twice every six months, this involves repetitive and time-intensive searching by the CRM. Thus, the most time-consuming part of the process is the search for a matching advocate, which is increased through the mentioned rejection loops.

One could argue that the simplest solution would be giving the SR the option to propose advocates that he or she considers a good fit. Before, this option existed, however, the SRs proposed the same most successful firms repeatedly. Thus, this option was disregarded. In the following section, a solution is proposed to optimize the identified problem area.

4.3. The Solution

The solution aims to reduce the number of loops taken to find a good advocate match for a request. This was achieved by enabling easier identification of disqualifying criteria of an advocate account or advocate persona. Furthermore, 53 variables instead of the mentioned three were used to determine the similarity between prospects and advocates. The proposed matches are automatically generated and made available for the CRM, which results in a less time-consuming process. This was accomplished by the implementation of a machine learning-driven algorithm that matches a request's associated account and sales opportunity data with already existent customer accounts and their closed and won sales opportunities. This resulted in two lists with best matches for a specific request: One with only those accounts that have shown a willingness for advocacy and the other one with all possible customer accounts. The second list was created to give the CRM the option to view those accounts and corresponding contacts that would be good advocates and should be contacted to join the customer reference program.

The lists are pushed into an application, called Reference Dashboard, which is hosted by the Customer Marketing Hub. This app displays an informed list of all requests that the CRM can click and see the 50 best matches among advocates and 50 best matches among all existing customers. Clicking on these matches enables the CRM to see corresponding contacts that are or are not available as advocates. Also, directly from the app, the advocate can be added to the request in the sales software, which triggers an email being sent out to the assigned CSM of the selected advocate. The following methodology section gives a deeper understanding of the machine learning model, the matching algorithm, and the creation of the Reference Dashboard.

5. Methodology

As mentioned above, a machine learning-driven approach was used to find the best matches for a specific request. Because it is found that a good match between accounts for reference purposes is based on their similarity, the model uses similarity metrics for matching. The company's focus lies on the optimization of certain use cases, which makes those a key variable for the model. Thus, in the first step, the model takes prospect account and opportunity data which are linked to the request object in the sales software to predict use cases the prospect company is likely to implement based on existing customers' implementations. The predicted use cases are then used for similarity matching of accounts in a second step.

First, the relevant data collected in the sales software was gathered and transformed. A data model (see Appendix H) that links all relevant tables via a 1:N pattern was created. These tables were then used in an analysis to gather, analyze, filter, and clean the data. The machine learning (ML) part was done using *JupyterLab*, an open-source interface that can contain data, live code, equations, visualizations, and text (Jupyter Project, n.d.). The data was pulled from the analysis to the *Jupyter Notebook* in the form of an *Online Analytical Processing Table* (OLAP) and used for further cleaning and analysis. The notebook was then used to create the model that finds the best matches for requests and each time it has run through, it pushes the results into the data model. From there, the results were embedded into the Reference Dashboard where the users can access the information. This process, from extraction to upload into the Reference Dashboard, is reloaded in an hourly manner to keep results up to date.

Step 1 – Use Case Prediction:

As each account can have multiple use cases, nine binary target features – one per use case – were created. With nine labels, the prediction of use cases was achieved using multilabel classification to predict one or more use cases for each prospect account. To find a suitable model for the multilabel classification, three models were compared and the best performing one was selected for account matching: *K-Nearest-Neighbors* (KNN), *Multi-layer Perceptron classifier* (MLP), and

Binary Relevance Gradient Boosting Classifier (BR GBM), which are explained in section 7. Randomized search cross-validation (CV) was used for hyperparameter tuning. To balance test and training sets for highly unbalanced labels, stratified sampling was used. Several metrics were employed to identify the strongest model: the *accuracy classification score*, the *jaccard score*, *precision* and *recall scores*, and the *f1-score*. The accuracy classification score returns the fraction of correctly classified samples, as shown in equation 1 (Pedregosa et al. 2011).

$$accuracy(y, \hat{y}) = \frac{1}{n_{samples}} \sum_{i=0}^{n_{samples}-1} 1(\hat{y}_i = y_i) \tag{1}$$

The jaccard score returns intersection size divided by the union size of the predicted labels and true labels, as shown in equation 2 (Pedregosa et al. 2011). When the average parameter is set to 'None', it returns the score for each label. Because the predicted labels are unbalanced, a weighted average of the score was computed. This is done by taking each label's score and computing their average weighted by support (Pedregosa et al. 2011).

$$jaccard(y_i, \hat{y}_i) = \frac{|y_i \cap \hat{y}_i|}{y_i \cup \hat{y}_i} \tag{2}$$

The precision score (equation 3) "is the ability of the classifier not to label as positive a sample that is negative" (Pedregosa et al. 2011), while the recall (equation 4) shows the number of truly predicted positives out of all positive samples (Pedregosa et al. 2011). The f1-score is a "harmonic mean of the precision and recall" that is computed using equation 5 (Pedregosa et al. 2011). All scores are calculated for each label as well as a weighted average.

$$precision = \frac{True\ Positives}{True\ Positives + False\ Positives} \tag{3}$$

$$recall = \frac{True\ Positives}{True\ Positives + False\ Negatives} \tag{4}$$

$$f1 = 2 * \frac{precision * recall}{precision + recall} \tag{5}$$

The model with the best scores was selected to predict the prospect's use cases. These newly won features were then added to the prospect's features to proceed with the second step.

Step 2: Similarity Matching

Using the *Nearest Neighbor* (NN) algorithm, the 50 best matches for one specific request were found, based on account, opportunity, and predicted use case data. This was done for two different lists: One that excludes all accounts that do not show a willingness for advocacy, and one that shows the similarity among all accounts. Both must fulfill certain criteria that identify the most successful accounts and are elaborated in the data preparation section. The NN model throws out the indices of these 50 best matches alongside their distance to the prospect data. The smaller the distance, the better the match. Thus, the distance can be used to sort the matches in an ascending manner in the resulting list. Having identified the best matching accounts, the results are hourly pushed into the Reference Dashboard, where they can be used by the CRM to identify reference matches.

6. Data Preparation

In the sales software used, reference requests are linked to an opportunity that is linked to an account and use case, just like existing customer accounts are linked to an opportunity and a use case. Opportunities are created for each sales option and use cases give further details, like use case or initiative information. Thus, one account can have many opportunities which can have many use cases. The use case data is aggregated on an account level, because, for this purpose, knowing use case details of the account induces more value than for each opportunity. However, because a request is primarily linked to an opportunity, the model considers multiple opportunities per account. Thus, the resulting data frame has one account which has one row per opportunity.

For the model, only accounts with an ARR of more than $0 and record type ‘Customer’ were considered, which excludes test or non-customer accounts. Only completed and won opportunities were selected, which excludes open or unsuccessful ones. This results in a data frame with X (censored due to data privacy) rows, 44 features, and 9 target features to predict the nine use cases of the prospect account. In the second step, the matching of accounts using NN, these nine use cases are added to the prospect data as additional features resulting in 53 features for the matching step.

Because the model aims at recommending only successful accounts, the data frames used for matching were filtered further. Risky accounts and accounts that have a subscription renewal coming up are excluded. The resulting number of accounts can quickly change if a filter does not apply anymore, or new accounts and opportunities are added.

In table 1, features extracted from the account and opportunity object are listed, where two are highlighted in yellow: Use case initiatives and objectives. These were, just like the predicted use cases, aggregated on an account level, which means that they were split into one boolean column per category. Also, the number of employees was categorized, as the CRM is using these categories for matching. The number of snap users, revenues, ACV targets, and expected opportunity revenue were split into bins to distribute the number of associated accounts more evenly.

Table 1: Features extracted from Account and Opportunity Object in the sales software (censored because of data privacy issues)

Because new accounts and opportunities are added every day, the following section shows a snapshot of the current situation. Table 2 lists target features that were used for predicting use cases, which are nine more general use case classifications, that the prospect is likely to implement. A more detailed use case classification was used to map higher level use cases that were not filled in the sales software (see Appendix F). The higher level use cases before and after this mapping are shown in figures 3 and 4. Figure 4 has more entries, which happens because a single use case might be mapped to another name than the name entered in the sales software. In this case, the account is assigned to both use case names.

Column Name	Types	Source	Description	Distribution	Distinct	Missing	Zeros
Account use cases	Categorical	Use Case	Use case aggregated on account level --> split into 9 columns	See below	9 cols	0%	0%
Use case 1	Boolean	Use Case		False: 476; True: 105	2	0%	0%
Use case 2	Boolean	Use Case		False: 569; True: 12	2	0%	0%
Use case 3	Boolean	Use Case		False: 327; True: 254	2	0%	0%
Use case 4	Boolean	Use Case		False: 351; True: 230	2	0%	0%
Use case 5	Boolean	Use Case		False: 473; True: 108	2	0%	0%
Use case 6	Boolean	Use Case		False: 553; True: 28	2	0%	0%
Use case 7	Boolean	Use Case		False: 519; True: 62	2	0%	0%
Use case 8	Boolean	Use Case		False: 570; True: 11	2	0%	0%
Use case 9	Boolean	Use Case		False: 390; True: 191	2	0%	0%

Table 2: Target Features extracted from Use Case Object in the sales software and aggregated on Account level (Status 12.12.2021) (Names and descriptions censored due to data privacy issues)

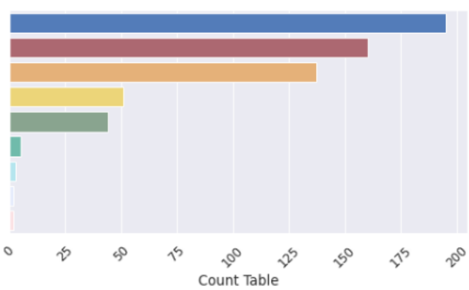

Figure 3: Use cases (before cleaning) (Status: 12.12.2021) (names censored due to data privacy issues)

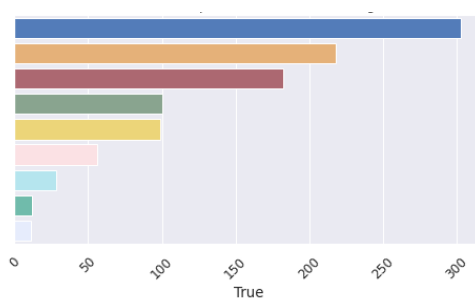

Figure 4: Use cases (after cleaning) (Status: 12.12.2021) (names censored due to data privacy issues)

The summary statistics of numerical features are shown in table 3. For all three the 75th percentile is close to the mean and thus higher numbers are rare. On average an account has 18 opportunities and almost 3 quotes.

Table 3: Summary statistics for numerical features (Status: 12.12.2021) (censored due to data privacy issues)

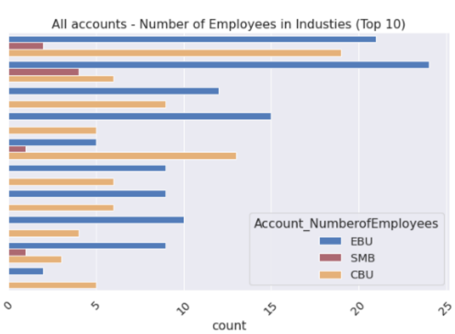

Figure 5: *All Accounts – Number of Employees in Industries*

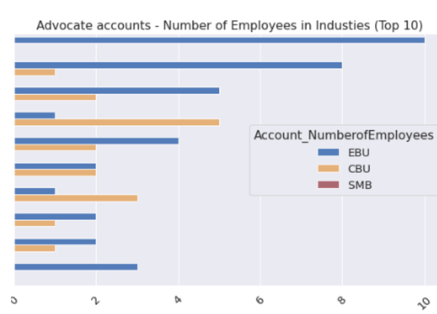

Figure 6: Advocate Accounts – Number of Employees in

(Top 10) (Status: 12.12.2021) (industry names censored due to data privacy issues)

Industries (Top 10) (Status: 12.12.2021) (industry names censored due to data privacy issues)

Figure 5 shows the top 10 industries of accounts, considering all customer accounts and their number of employees within categories. Only a small number of small and midsize businesses (SMB) is available.

Considering all accounts, four regions are represented, namely Europe, North America (NA), Asia-Pacific, and Latin America, while only Europe and NA are present once the list is filtered to only advocate accounts, as shown in figures 7 and 8.

Figure 7: All Accounts GFTS accounts per region (Status: 12.12.2021 (censored due to data privacy issues)

Figure 8: Advocate Accounts – GFTS accounts per region (Status: 12.12.2021) (censored due to data privacy issues)

7. Machine Learning Model Selection

As mentioned in the methodology section, three models were compared and the best one chosen to predict use cases for the prospect account: KNN, MLP, and BR GBM. Because nine use case columns were to be predicted, the models used are compatible with multi-label classification. The best parameters for the models are found using randomized search CV. This section describes how the different models work and how they performed in the prediction. It must be noted that because data changes frequently, the numbers mentioned are only a snapshot of the current situation.

K-Nearest-Neighbors (KNN): KNN is a supervised learning algorithm that measures the distance from prospect features to considered existing accounts within a matrix. The K closest accounts are selected, and a larger vote wins for predicting the target features. KNN is a lazy learner, which "implies that it does not learn a discriminative function from the training data but rather memorizes the training data instead" (Vatsal 2021). It is flexible in using different distance metrics (Vatsal 2021), however, it performs worse on imbalanced datasets and the larger the data, the slower its computational speed (Vatsal 2021). For KNN the following best parameters were identified:

Parameter	Description (Pedregosa et al. 2011)	Value
Number of neighbors (K)	Number of neighbors	2
Weights	Weight function used: 'uniform': All points weighted equally; 'distance': Closer neighbors have greater influence	Distance
p	p = 1: Manhattan Distance; p = 2: Euclidean Distance	2

Table 4: KNN Parameters, Descriptions from Scikit Learn (Pedregosa et al. 2011) (Status: 12.12.2021)

Multi-layer Perceptron Classifier (MLP): The MLP is a supervised neural network model that aims at minimizing the loss function (Pedregosa et al. 2011). It has a minimum of three layers: the input layer, one or more hidden layers, and an output layer, where "the data flows in the forward direction from input to output layer" (Raj and Lin 2020). The model uses backpropagation, which "repeatedly adjusts the weights of the connections in the network so as to minimize a measure of the difference between the actual output vector of the net and the desired output vector" (Kostadinov 2019). MLP can be applied to non-linear problems and has a high computational power with large datasets (Sumsion et al. 2019). Parameters found for MLP are the following:

Parameter	Description (Pedregosa et al. 2011)	Value
Activation function	Activation function for hidden layer: 'identity': no-op activation, returns f(x) = x; 'logistic' logistic sigmoid function, returns f(x) = 1/(1+exp(-x)); 'tanh': hyperbolic tan function, returns f(x) = tanh(x); 'relu': rectified linear unit function, returns f(x) = max(0,x)	relu
Solver	Solver for weight optimization: 'lbfgs': optimizer in the family of quasi-Newton methods; 'sgd': stochastic gradient descent; 'adam': stochastic gradient-based optimizer	sgd
Alpha	L2 penalty (regularization term) parameter	0.001
Learning rate	Learning rate schedule for weight updates (only used when solver = 'sgd'): 'constant': constant learning rate given by 'learning_rate_init'; 'invscaling': learning_rate = learning_rate_init / pow(t,power_t); 'adaptive': keeps the learning rate constant to 'learning_rate_init' if training loss keeps decreasing.	Invscaling
Learning rate init	Initial learning rate used. Controls the step size in updating the weights. Only used when solver='sgd' or 'adam'.	0.01
Power T	Only used when solver='sgd': The exponent for inverse scaling learning rate.	0.1
Max iter	Maximum number of iterations. For stochastic solvers ('sgd', 'adam'), note that this determines the number of epochs, not the number of gradient steps.	200
Early stopping	If set to True, terminates training early when validation score is not improving further.	False

Table 5: MLP Parameters, Descriptions from Scikit Learn (Pedregosa et al. 2011) (Status: 12.12.2021)

Binary Relevance Gradient Boosting Classifier (BR GBM): Gradient booster models (GBM) "train upon the residuals, the difference between the prediction and the actual results" (Wade 2020). Thus, for each iteration, it considers previous errors and improves upon this, which aims to turn weak learners, like decision trees, into strong learners. GBM is flexible in dealing with different loss functions, and imputation of missing data is not required.(Kurama 2019). Nonetheless, GBM continues until all errors are minimized, which can "overemphasize outliers and cause overfitting" (Kurama 2019). BR classification makes the GB classifier compatible with multiple labels. The BR algorithm "transforms a multi-label classification problem with L labels into L single-label separate binary classification problems using the same base classifier provided in the constructor" (Szymanski and Kajdanowicz 2017) This means that it uses all given features to predict each use case separately. This approach is computationally efficient; however, it ignores a possible correlation between target features (Hong 2014). The parameters used for GBM are the following:

Parameter	**Description** (Pedregosa et al. 2011)	**Value**
Maximum Depth	Maximum depth of the individual regression estimators.	4
Learning Rate	Shrinks the contribution of each tree by the value.	0.4
Maximum Features	The number of features to consider when looking for the best split: ‘auto’ & ‘sqrt’ = max_features = sqrt(n_features); ‘log2’ = max_features = log2(n_features); None = max_features = n_features	Auto

Table 6: Parameters BR GBM Classifier, Descriptions from Scikit Learn (Pedregosa et al. 2011) (Status: 12.12.2021)

7.1.Model Performance

Model	**Overall Accuracy**	**Weighted Jaccard Score**	**Weighted Precision Score**	**Weighted Recall Score**	**Weighted F1 Score**
K Nearest Neighbors (KNN)	42.37%	52.13%	72.07%	66.67%	67.65%
Multilayer Perceptron Classifier (MLP)	35.59%	53.76%	69.17%	67.57%	67.78%
Binary Relevance Gradient Boosting (BR GBM)	45.76%	58.92%	83.86%	65.77%	72.86%

Table 7: Model Performances (Status: 12.12.2021)

In table 7, the weighted scores, as well as the overall accuracy of the three models, are compared. The weighted average is taken because the predicted labels are very unbalanced and perform very differently. This is also why the accuracy score, which is an unweighted average of the results, scores worse than the weighted jaccard score. It becomes evident, that, looking at the scores in table 7, the binary relevance gradient boosting model performs best in all scores. To get a better understanding of how the scores are calculated, the following table shows the scores for the individual use case predictions for the best performing model: Binary Relevance Gradient Boosting.

Metric	**Use Case 1**	**Use Case 2**	**Use Case 3**	**Use Case 4**	**Use Case 5**	**Use Case 6**	**Use Case 7**	**Use Case 8**	**Use Case 9**
Distribution in training	False: 476 True: 105	False: 569 True: 12	False: 327 True: 254	False: 351 True: 230	False: 473 True: 108	False: 553 True: 28	False: 519 True: 62	False: 570 True: 11	False: 390 True: 191
Jaccard Score	54.54%	100%	58.75%	61.54%	66.67%	20%	37.5%	50%	41.67%
F1 Score	70.59%	100%	85.71%	76.19%	80%	33.33%	54.55%	66.67%	58.82%
Prediction	85.71%	100%	90%	84.21%	88.89%	50%	100%	100%	66.67%
Recall	60%	100%	81.82%	69.57%	72.73%	25%	37.5%	50%	52.63%

Table 8: Binary Relevance Gradient Boosting Performance (Status: 12.12.2021 (use case names censored due to data privacy)

In this snapshot of model performance in table 8, use case 2 stands out, scoring 100% for all metrics.

This means that all predictions for this use case are correct (see Appendix G). If the use case distribution is very balanced or very unbalanced between the binary values, they usually score better than the ones in between. The use case prediction made using the BR GBM were added as additional features to the prospect's data and used to match the prospect with existing accounts.

8. Matching based on similarity

For the matching process, which is based on the similarity of features between a prospect and existing customer accounts, the unsupervised NN algorithm from scikit learn was used. The algorithm finds a defined number of nearest neighbors to the prospect features. For this model, the algorithm predicted the 50 nearest neighbors of the prospect. From this, a mapping list was created consisting of the reference request name, the matched account ID, and the distance from the account to the prospect determined by the NN algorithm. Thus, for each prospect request, 50 entries were created in this table and ordered in ascending order of distance. Because an account can have several rows dependent on the number of opportunities, only the row with the closest distance to the prospect data was kept. This was done once for only advocate accounts and once for all accounts resulting in two mapping tables. The mapping tables are then pushed into the data model (see Appendix H), from where the data gets called into the Reference Dashboard. This is then used by the CRM to find good matches for a request.

9. Result Presentation & Measure of Success

The Reference Dashboard stores all requests with the corresponding prospect information in a list (see Appendix I). The entries are ordered by creation date of the request to have the most recent ones at the top. The list can be filtered for non-completed requests and one can click on a request to open a pop-up with a detailed view of this request, including the two lists of proposed matches (see Appendix J). The first list shows matches among advocate accounts and the second considers all existing customers. The matched accounts are clickable to open another window (see Appendix K) with details on account contacts, where an advocate can be requested via a button.

To improve performance, the models will be closely monitored, and parameters adjusted. The success of the Reference Dashboard will be measured by looking at the adoption metrics, such as daily active users, which can be tracked within the app. If the dashboard was not useful to the CMD, it is assumed that it would not be used. Feedback from the CMD, who will be the power users of the dashboard, will be considered to improve the ease of use of the Reference Dashboard. Furthermore, the time taken from request creation to the reference meeting taking place will be assessed and compared to the previous manual process in the form of A/B testing.

10. Limitations & Future Work

Due to time and publishing constraints, the dashboard and therefore the model is not actively used yet at the company. Publishing will most likely take place within the coming two months as soon as the CMH has advanced further. Therefore, the measurement of success of the model and dashboard could not be executed yet but will be performed as soon as the model is published to users.

Because, applying the filters, currently only a limited amount of advocate companies is available for advocacy, oftentimes same accounts are recommended at top rows. Also, in case the CRM has a more specific search, such as a certain industry, only some accounts are available for each. The more accounts engage in advocacy, the better the model will work and find good matches among the existing accounts. This limitation should, in the future, be improved by increasing account engagement in co-marketing activities. Furthermore, because some use cases are more frequently implemented than others, their performance is much better within the model. By expanding on opportunities at existing customers, this distribution will likely improve over time.

Because data on references and requests is limited and unreliable, it cannot be used as a feedback form to the recommendations. With the CMH, the tracking of advocacy will become much easier in the future which can also be used to make the Reference Dashboard more comprehensive. For example, a new feedback form with more detailed questions is created to gather the sentiment of the prospect after the reference meeting has taken place. This data can be used to train the model whether a predicted match was a good one or not to improve the recommendation system.

11. Conclusion

The importance of ABM and customer engagement becomes evident in B2B companies where 20% of accounts generate about 80% of the revenue. Accounts must be well managed from beginning to end of the sales cycle. The customer sales reference process enables this engagement at the very start of the account's journey. Being well-matched with an existing account and learning about its experiences, will strengthen the prospect's trust and feeling of security towards B2B software companies.

The proposed model helps to improve the matching process by automatically recommending the best matches for a specific prospect and making this information easily accessible for the CRM. The model is created in a two-step process: First, the prospect's most probable use cases for implementation in the software are predicted using Multilabel Classification. Among three models, the gradient boosting classifier model was identified as the best performing one and was made compatible for multi-label classification by using the binary relevance algorithm. In a second step, using this prediction as an additional feature, the prospect account is matched with existing accounts based on similarity.

Due to partially unbalanced target features, namely use cases, some use case predictions score much higher than others. With more and more customers being acquired and more and more opportunities being created, the diversity of use cases implemented will most probably become stronger, which will deliver a more balanced dataset in the long term and therefore stronger models.

II. References

Aarikka-Stenroos, Leena, and Elina Jaakkola. 2012. "Value Co-Creation in Knowledge Intensive Business Services: A Dyadic Perspective on the Joint Problem Solving Process." *Industrial Marketing Management* 41 (1): 15–26. https://doi.org/10.1016/J.INDMARMAN.2011.11.008.

Bacon, Andy. 2020. "Your Essential Guide to Proving Success in Account Based Marketing | LinkedIn." February 7, 2020. https://www.linkedin.com/pulse/your-essential-guide-proving-success-account-based-andy-bacon-f-idm/.

———. 2021. *Strategic Account-Based Marketing: How to Tame This Beast*. Springer, Cham. https://doi.org/10.1007/978-3-030-54292-4_17.

Bijmolt, Tammo H.A., Peter S.H. Leeflang, Frank Block, Maik Eisenbeiss, Bruce G.S. Hardie, Aurélie Lemmens, and Peter Saffert. 2010. "Analytics for Customer Engagement." *Journal of Service Research* 13 (3): 341–56. https://doi.org/10.1177/1094670510375603.

Blazevic, Vera, Wafa Hammedi, Ina Garnefeld, Roland T. Rust, Timothy Keiningham, Tor W. Andreassen, Naveen Donthu, and Walter Carl. 2013. "Beyond Traditional Word-of-Mouth: An Expanded Model of Customer-Driven Influence." *Journal of Service Management* 24 (3): 294–313. https://doi.org/10.1108/09564231311327003.

Burgess, Bev. 2016. "Driving B2B Growth with Account-Based Marketing." *Market Leader*. www.itsma.com.

Field Marketing Department. 2021. "Account Based Strategy."

Connelly, Brian L., S. Trevis Certo, R. Duane Ireland, and Christopher R. Reutzel. 2011. "Signaling Theory: A Review and Assessment." *Journal of Management* 37 (1): 39–67.

https://doi.org/10.1177/0149206310388419.

Gene Day, Daniel, and Savannah Wei Shi. 2020. “Automated and Scalable: Account-Based B2B Marketing for Startup Companies.” *Journal of Business Theory and Practice* 8 (2): p16. https://doi.org/10.22158/JBTP.V8N2P16.

Golec, Chris, Peter Isaacson, and Jessica Fewless. 2019. *Account-Based Marketing: How to Target and Engage the Companies That Will Grow Your Revenue*. New Jersey: John Wiley & Sons. Inc. Hoboken. https://books.google.de/books?hl=de&lr=&id=QNOFDwAAQBAJ&oi=fnd&pg=PA1&dq=account-based+marketing&ots=d3qpeS4zPu&sig=BW1QQFLwtNi8CF15RE3ugBan9e8&redir_esc=y#v=onepage&q=account-based%20marketing&f=false.

Harmeling, Colleen, Jordan Moffett, and Brad D Carlson. 2016. “Toward a Theory of Customer Engagement Marketing.” *Article in Journal of the Academy of Marketing Science*. https://doi.org/10.1007/s11747-016-0509-2.

Hong, Charmgil. 2014. “Multi-Label Classification.”

Jaakkola, E., and L. Aarikka-Stenroos. 2019. “Customer Referencing as Business Actor Engagement Behavior – Creating Value in and beyond Triadic Settings.” *Industrial Marketing Management* 80 (July): 27–42. https://doi.org/10.1016/J.INDMARMAN.2018.06.014.

Jaakkola, Elina, Leena Aarikka-Stenroos, and Allan J. Kimmel. 2014. *Customer Experience Management: Enhancing Experience and Value through Service Management*. https://www.researchgate.net/publication/271512730_Leveraging_customer_experience_communication.

Jaakkola, Elina, and Matthew Alexander. 2014. “The Role of Customer Engagement Behavior in

Value Co-Creation: A Service System Perspective." *Journal of Service Research* 17 (3): 247–61. https://doi.org/10.1177/1094670514529187.

Jupyter Project. n.d. "Https://Jupyter.Org/."

Kostadinov, Simeon. 2019. "Understanding Backpropagation Algorithm." *Towards Data Science*, August. https://towardsdatascience.com/understanding-backpropagation-algorithm-7bb3aa2f95fd.

Kruse, Kevin. 2016. "The 80/20 Rule And How It Can Change Your Life," March 7, 2016. https://www.forbes.com/sites/kevinkruse/2016/03/07/80-20-rule/?sh=3df3d4273814.

Kurama, Vihar. 2019. "Gradient Boosting in Classification: Not a Black Box Anymore!" *Paperspace Blog*. https://blog.paperspace.com/gradient-boosting-for-classification/.

Malthouse, Edward C., Michael Haenlein, Bernd Skiera, Egbert Wege, and Michael Zhang. 2013. "Managing Customer Relationships in the Social Media Era: Introducing the Social CRM House." *Journal of Interactive Marketing* 27 (4): 270–80. https://doi.org/10.1016/J.INTMAR.2013.09.008.

Morgado, Andre Vilares, Luis Mota, and D E Castro. 2015. "CUSTOMER REFERENCING: TESTING A CONCEPTUAL FRAMEWORK IN THE CONTEXT OF A CASE STUDY."

Pavan Kumar, G, and K Rajasekhar. 2020. "Account Based Marketing in B2B Industry." *Journal of Interdisciplinary Cycle Research*.

Pedregosa, F, G Varoquaux, A Gramfort, V Michel, B Thirion, O Grisel, M Blondel, et al. 2011. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research* 12: 2825–30.

Povolna, Lucie. 2017. "MARKETING COMMUNICATIONS ON B2B MARKETS."

https://doi.org/10.7441/dokbat.2017.29.

Raj, Pethuru, and Jenn Wei Lin. 2020. “Stepping into the Digitally Instrumented and Interconnected Era.” *Advances in Computers* 117 (1): 1–34. https://doi.org/10.1016/bs.adcom.2019.09.008.

Salminen, Risto T., and Kristian Möller. 2006. “Role of References in Business Marketing-towards a Normative Theory of Referencing.” *Journal of Business-to-Business Marketing* 13 (1): 1–51. https://doi.org/10.1300/J033V13N01_01.

Stadelmann, Martin, Mario Pufahl, and David D. Laux. 2020. *CRM Goes Digital: Digitale Kundenschnittstellen in Marketing, Vertrieb Und Service Exzellent Gestalten Und Nutzen*. Sales Excellence. Springer Gabler. https://books.google.de/books?hl=de&lr=&id=pwTaDwAAQBAJ&oi=fnd&pg=PA160&dq=account-based+marketing&ots=lFmIh5GoKZ&sig=VeZpuAEraws05NepE5TG2lglnRc&redir_esc=y#v=onepage&q=account-based%20marketing&f=false.

Sumsion, G. Rex, Michael S. Bradshaw, Kimball T. Hill, Lucas D.G. Pinto, and Stephen R. Piccolo. 2019. “Remote Sensing Tree Classification with a Multilayer Perceptron.” *PeerJ* 2019 (2). https://doi.org/10.7717/PEERJ.6101/FIG-4.

Szymanski, P., and T. Kajdanowicz. 2017. “Scikit-Multilearn: Multi-Label Classification in Python — Multi-Label Classification for Python.” February 2017. http://scikit.ml/api/skmultilearn.html.

Vajre, Sangram. 2016. *Account-Based Marketing For Dummies*. https://books.google.de/books?hl=de&lr=&id=Rw79CgAAQBAJ&oi=fnd&pg=PR17&dq=account-

based+marketing&ots=QRyqqh3th7&sig=vqOHdc2I253FDMqYzAytaDvBgLA&redir_esc=y#v=onepage&q=account-based%20marketing&f=false.

Vatsal. 2021. "K Nearest Neighbours Explained." *Towards Data Science*, January 11, 2021. https://towardsdatascience.com/k-nearest-neighbours-explained-7c49853633b6.

Wade, Corey. 2020. "Getting Started with XGBoost in Scikit-Learn." *Towards Data Science*, November 10, 2020. https://towardsdatascience.com/getting-started-with-xgboost-in-scikit-learn-f69f5f470a97.

Zhang, J. 2018. "Product Market Effects of Customer Referencing." https://www.microsoft.com/hk/casestudies/.

Name censored 2021. "Interview with Director of Field Marketing."

Name censored. 2021."Interview with Senior Field Marketing Manager."

III. Appendices

Appendix A: Comparison of Inbound-Marketing and Account-based Marketing (Stadelmann, Pufahl, and Laux 2020)

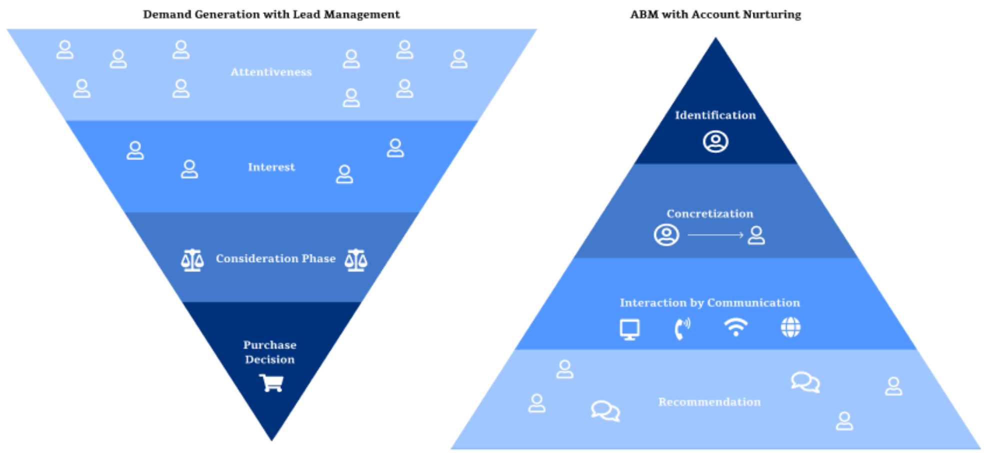

Appendix B: The three types of account-based marketing (Pavan Kumar and Rajasekhar 2020; Golec, Isaacson, and Fewless 2019)

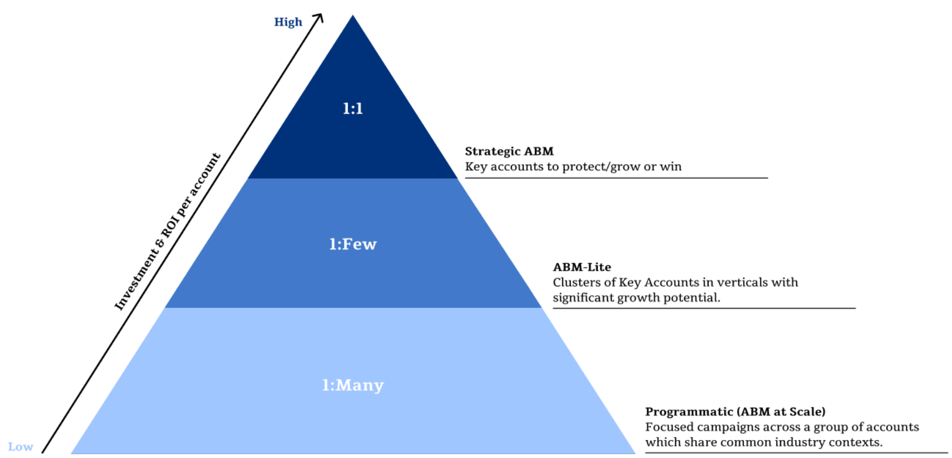

Appendix C: Interview with Director of Field Marketing of the software company looked at (20.10.2021). This interview is taken out due to data privacy reasons of that company.

Appendix D: Interview with Senior Field Marketing Manager (21.10.2021). This interview is taken out due to data privacy reasons of that company.

Appendix E: Signaling Timeline in Company Context (Connelly et al. 2011)

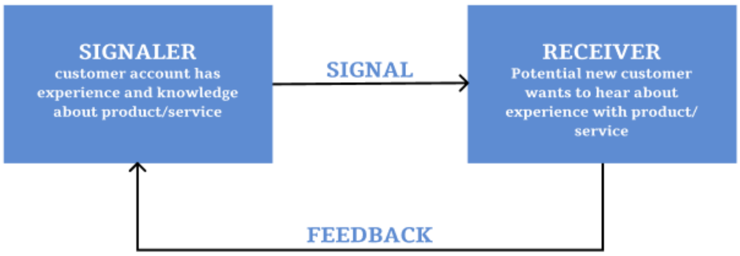

Appendix F: Use Case Mapping for Target Features. This mapping table is taken out due to data privacy reasons.

Appendix G: Confusion matrices for use case prediction (use case names censored due to data privacy issues)

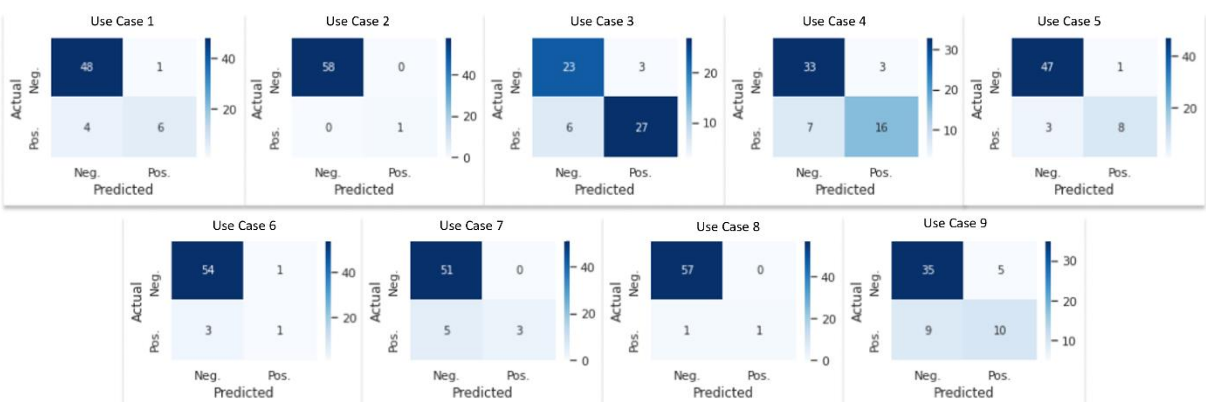

Appendix H: Data Model

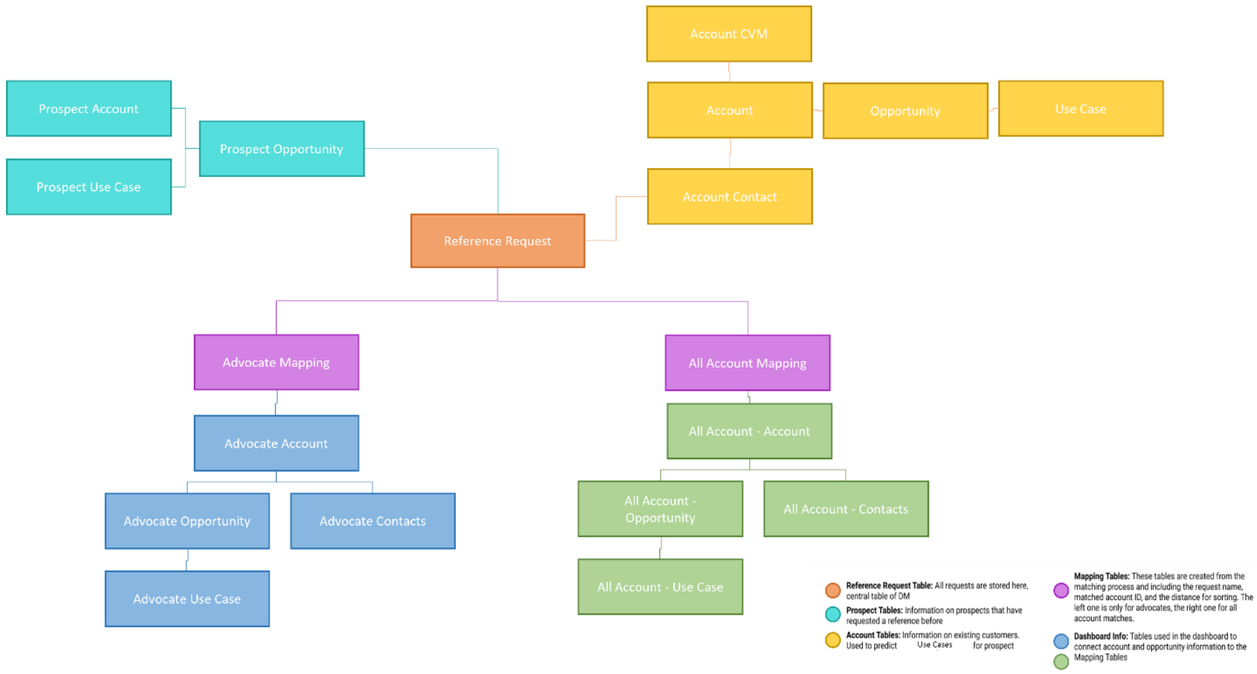

Appendix I: Reference Dashboard – List of Requests. This view is taken out due to data privacy reasons of that company.

Appendix J: Reference Dashboard – Detail View of Requests with proposed best matches. This view is taken out due to data privacy reasons of that company.

Appendix K: Reference Dashboard – Detail View of matched accounts with contact information This view is taken out due to data privacy reasons of that company.