

2020

Predicting Attrition - a Driver for Creating Value, Realizing Strategy, and Refining Key HR Processes

Kevin Mendonsa

Southern Methodist University, kmendonsa@mail.smu.edu

Maureen Stolberg

Southern Methodist University, mstolberg@smu.edu

Vivek Viswanathan

Southern Methodist University, vviswanathan@mail.smu.edu

Scott Crum

MSCI Inc., Scott.Crum@msci.com

Follow this and additional works at: <https://scholar.smu.edu/datasciencereview>



Part of the [Business Administration, Management, and Operations Commons](#), [Data Science Commons](#), [Human Resources Management Commons](#), [Management Sciences and Quantitative Methods Commons](#), and the [Strategic Management Policy Commons](#)

Recommended Citation

Mendonsa, Kevin; Stolberg, Maureen; Viswanathan, Vivek; and Crum, Scott (2020) "Predicting Attrition - a Driver for Creating Value, Realizing Strategy, and Refining Key HR Processes," *SMU Data Science Review*. Vol. 3: No. 2, Article 2.

Available at: <https://scholar.smu.edu/datasciencereview/vol3/iss2/2>

This Article is brought to you for free and open access by SMU Scholar. It has been accepted for inclusion in SMU Data Science Review by an authorized administrator of SMU Scholar. For more information, please visit <http://digitalrepository.smu.edu>.

Predicting Attrition - a driver for Creating Value, Realizing Strategy, and Refining Key HR Processes

Kevin Mendonsa¹, Maureen Stolberg¹, Vivek Viswanathan¹, and Scott Crum²

¹ Master of Science in Data Science, Southern Methodist University, Dallas TX
75275 USA {kmendonsa,mstolberg,vviswanathan}@smu.edu

² Chief Human Resources Officer, MSCI Inc., 7 WTC, Greenwich Street. NY 10007

Abstract. Talent is the most important asset for every organization’s success. While attrition (or churn) and turnover can refer to both employees and customers, this paper will focus on employee attrition only. Many organizations accept attrition as an inevitable cost of doing business and do nothing to adopt or implement mitigating strategies to combat it. World class companies on the other hand take deliberate measures to understand, control and mitigate attrition (turnover) at every stage. Unmitigated attrition can have a devastating effect on an organization’s bottom line and market value. In addition, the “invisible” costs of low employee morale, reduced employee engagement, stagnant innovation are more harmful to the well-being of any organization. Predicting employee attrition allows organizations reasonable time to have discussions with employees predicted to leave, in order to retain them if aligned with strategy. It also enables the organization to develop alternatives to proactively address attrition by building appropriate talent pipelines and conducting loss impact analyses especially for key roles and strategic projects. The aim of this paper is to highlight the importance of talent to the success of the organization, its impact to profits and overall market value. We intend to provide a framework for data collection methodologies and the prediction of employee attrition by analyzing multiple factors and attributes using defined machine learning classification techniques and models.

Keywords — employee attrition, turnover, predictive analytics, data mining, churn prediction, machine learning

1 Introduction

Talent is the backbone of the organization and represents 70% of the costs in most organizations [11]. Per compensation data surveys, over the past five years, total turnover has increased from 15.1% to 18.5% [19]. Absent good talent, an organization cannot survive or compete successfully. Bill Gates was once quoted as saying, “You take away our top 20 employees and we [Microsoft] become a mediocre company” [19]. Jim Goodnight, Co-Founder & Chief Executive Officer

of SAS admits “Creativity is especially important to SAS because software is a product of the mind. Ninety-five percent of my assets drive out the gate every evening. It’s my job to maintain a work environment that keeps those people coming back every morning”¹. While talent is important to all organizations it is even more critical to technology, knowledge based and service organizations. These organizations are 100% dependent on their employees to deliver innovations, breakthroughs and services that create value. In a knowledge driven economy, talent being the key competitive differentiator, retaining it is of paramount importance [7]. While talent acquisition itself is difficult amid the current “war for talent”, it is not adequate in today’s changing economic landscape; employee retention is more critical. Losing skilled, knowledgeable and well-trained employees can seriously impact an organization’s productivity, innovation, performance and shareholder value in the market [7]. According to Harvard Business Review article “How to Keep your Top Talent”, one-quarter of the highest potential people in any organization intend to jump ship within the year [12]. Emerging leaders require stimulating work, lots of recognition, and the chance to prosper, else they can quickly become disenchanted and disengaged. Managing the quantity and quality of high potentials, especially at the corporate level, are extremely important versus delegating their development to line managers, which could further limit their opportunities. Another factor that determines a rising star’s engagement is the sense of recognition, which is primarily through pay. “A” players must be offered differentiated compensation and recognition. “In a September 2009 survey by the Corporate Executive Board, one in three emerging stars reported feeling disengaged from his or her company. Even more striking, 12% of all the high potentials in the organizations studied said they were actively searching for a new job. This suggests that as the economy rebounds and the labor market warms up, organizations may see their most promising employees take flight in large numbers” [12]. The disenchantment among high potentials can have troubling consequences for many companies. The discretionary effort (the crucial willingness to go above and beyond) can be as much as 50% lower among highly disengaged employees than with employees with average engagement. No CEO or organization can afford to lose so much productivity from its core talent and high potentials. Keeping young stars engaged is an even greater challenge for senior management who must double or even triple its efforts to accomplish this. Claudio Fernández-Aráoz, Boris Groysberg, and Nitin Nohria in their article on “Emerging best practices in managing your company’s future leaders” and “How to Hang on to your High Potentials” proffer that the war for talent shows no signs of letting up, even in sectors experiencing modest growth. “According to a global study conducted, only 15% of companies in North America and Asia believe that they have enough qualified successors for key positions. The picture is slightly better in Europe, but even so, fewer than 30% of European companies feel confident about the quality and amount of talent in their pipelines” [4]. Moreover, in emerging markets where many companies are focus-

¹ More information may be found at https://www.sas.com/en_us/company-information/leadership/jim-goodnight.html. Last accessed 22 Mar, 2020.

ing their growth strategies, the supply of experienced managers are the most limited, and the shortage is expected to continue for another two decades.

Attrition and turnover are key performance metrics for any organization. Understanding the drivers of attrition are crucial to address underlying organizational issues and develop remediating strategies. Employee turnover is typically split into two types:

1. Voluntary turnover or Attrition is initiated by the employee, for example, a worker quits to take another job. World class organizations further classify and categorize voluntary attrition as:
 - (a) Regrettable Attrition: This comprises Top Talent or High Performing employees also referred to as Dysfunctional turnover and is much harder to replace. This can have an even more pronounced impact when they join a competitor. It could also be employees with hard to find skills or departures of women or minority group members that erode the diversity of an organization's workforce [11]. Regrettable attrition is further stratified as:
 - i. Top Talent/High Potential: These are roles that have the greatest potential within an organization to drive the strategy and can grow into other higher roles within the organization. They are usually considered for upward mobility or as successors to other roles.
 - ii. High Impact: These roles are also critically important and drive the most impact in the growth and strategy of an organization. They are typically roles with incumbents that have the most experience and knowledge within the organization.
 - (b) Non-Regrettable Attrition: It refers to roles that while important for the smooth functioning of any organization are also easier to fill. They are non-strategic or non "keyman" roles and are usually the first positions during a re-organization that could be impacted.
 - i. Contributors: These are roles that very necessary and are responsible for day-to-day operations or "keeping the lights on" and moving things forward. They range from administrative roles to other necessary roles to maintain daily production or operations.
2. Involuntary turnover is initiated by the organization, for instance, an organization terminates an employee for poor performance or restructuring. This is also referred to as Functional turnover and includes employees whose talents are easy to replace. Involuntary is also sometimes categorized as:
 - (a) For Cause – This could be for several reasons but primarily related to violations of an organization's Code of Conduct and other policies.
 - (b) Reorganizations
 - i. Restructuring – for improvements in productivity or relocating for labor arbitrage
 - ii. Cost cutting measures – Reducing the cost footprint for operating profits
 - iii. Divestitures – The sale of a division of an organization
 - iv. Mergers – Often there could be a duplication or redundancy of certain roles

The aim of this study is to provide a framework for using machine learning techniques to construct a solution-oriented model that uses the results of predicting employee attrition to provide recommendations for maximizing a retention campaign's ROI and strategically managing the talent pool. Our objective is two-fold. The first is to identify employees that are at risk of leaving and determine what are the factors driving this decision for each employee. The second is to maximize a retention campaign's ROI by identifying employees with the highest estimated cost of attrition so HR can make the most cost-effective choices when exercising retention strategies.

The rest of the paper is organized as follows: Section II, describes the organizational and economic costs of attrition and turnover and the imperative for predicting it so mitigating strategies can be employed. Section III identifies and defines the factors that typically impact attrition – both extrinsic and intrinsic to the organization. In Section IV, data collection methodologies, HRIS systems of record and exit surveys by 3rd party organizations are discussed along with the modeling approaches. Section V describes the various models and comparisons between them. In Section VI, the model effectiveness parameters are defined and compared. Section VII concludes all the findings of the work done in this paper [19]. Finally, Section VIII discusses ethical considerations.

2 Economic Costs of Attrition

In today's competitive business environment, the impact of attrition can be detrimental to both the bottom line and employee morale. To decrease attrition, managers must understand the root causes of why their most valuable assets are leaving. "Companies must prioritize investigating the causes of regrettable and non-regrettable attrition and build this into a rigorous HR process. Was it avoidable? Could the departure have been avoided by better management? Should hiring take more account of certain demographic factors which influence tenure?" [11] Are there demographic factors that influence tenure that should be considered during the hiring process? Does more careful attention need to be given to "invisible factors" such as a candidate's cultural fit. Greater focus is required to ensure that a candidate's job and role expectations align with the organization's expectations.

Attrition could occur for many different reasons and categorizing and understanding the root causes is crucial to addressing it. It impacts an organization's ability to deliver critical projects and drive the success of short and long-term strategies. It can be quantifiable in some instances but often also has hidden costs that are difficult to measure such as employee morale. It directly impacts the bottom line and erodes profits and shareholder value. Often the cost of filling an open position can be between 150% to 200% of the position's salary.

The costs of attrition range from quantifiable numbers to hidden costs. When employees resign from companies, costs are incurred in recruiting new employees and training them. Productivity will be lower until new hires learn the business. In a customer-based business, customers could potentially become dissatisfied

if the new hire is not proficient in the role. The business could lose customers who are dissatisfied with the service or because they miss the relationship they enjoyed with the previous employee. Revenue would decrease. While attrition is inevitable in every organization, understanding the organizational costs and wider economic costs are essential to stem losses and impacts to productivity. “According to our research, based on PwC figures, up to \$27B is being wasted in the US economy alone because people are not hiring the right candidates: clearly there is a serious issue here that needs to be addressed” [11]. First year attrition is particularly concerning for organizations, as research has shown that employees typically become productive contributing assets after their first year of employment. It takes time for employees to fit into the culture, and companies spend a lot of resources training new hires and assimilating them into the organization. First-year-of-service turnover among US organizations is around 24%, according to PwC², and each leaver costs the organization 1 to 1.2 times their annual salary [11] in addition to the cost already incurred to acquire them. “For example, in a hypothetical 10,000-employee organization with a 10% hiring rate, 240 of its 1,000 annual recruits are being lost. For an employee earning \$50,000, the costs of recruiting and retraining would be \$50,000 to \$60,000. In this hypothetical organization, the ultimate cost of attrition would be \$60,000x240 times over – a figure close to \$15M. If this organization could hang on to just 10% of the 240 employees leaving, they could potentially save more than \$1M annually. The cost impact to the wider economy are extraordinary even when using conservative estimates. It is estimated that on average 5.1 million Americans quit or were laid off from their jobs in December 2016 making it the biggest month for job separations since 2008. Using more conservative estimates, of the 1.5M voluntary resignations every month, only 10% are addressable. Assuming a very conservative cost of attrition at 30% of the average US salary still results in an estimated \$27B cost to the US economy” [11].

3 Factors of Attrition

Market Environment factors have also been shown to be associated with employee turnover, reflecting the perception of alternative employment opportunities as employees periodically evaluate their current situation (Abelson, 1986) [1].

This section examines how various factors in an organization’s business environment can influence employee attrition. For the purposes of this study, we will focus on three specific categories of influential factors. These are macro-environment factors, micro-environment factors, and internal environment factors.

² Additional information can be found on page 8 at <https://www.pwc.com/us/en/hr-management/publications/assets/pwc-trends-in-workforce-analytics.pdf>. Last accessed 22 Mar, 2020.

3.1 Macro-environment Factors

Macro-environment factors are external factors that influence business operational growth and long-term sustainability. These factors cannot be controlled but can be understood by identifying contributing environmental forces.

According to Francis J. Aguilar's 1964 publication titled, "Scanning the Business Environment" [2], the four broad types of macro-environmental factors and forces that are interrelated and affect organizations include but are not limited to: economic, technical, political, and sociocultural factors.

Economic Factors are the determinants that can influence a country's ability to manufacture or produce domestic goods outside normal commerce activities. Directional changes to the aggregate GDP create business cycles that impact a country's economic value in many different ways. During periods of economic growth, production output grows, consumer spend, the economy expands causing income levels to rise while keeping unemployment rates relatively low. Factors associated with supply and demand affect how prices, labor, and quantity of goods impact free market trade activity.

Two types of economic factors that influence the labor market are broad aggregate factors and market specific factors. They directly influence labor market related factors that impact the supply and demand trends. Broad aggregate factors consider various aspects of the country's economy, such as changes in gross domestic production output (GDP), interest rates, consumer prices, and the disposition of personal income.

Market specific factors that directly relate to changes in labor market supply & demand trends include changes in labor participation, employee turnover (labor supply), job availability (labor demand), workforce flows (labor flows), part-time contractor labor demand, and work compensation trends, that are comprised of salary, benefits, and wage differentials.

Technological Factors are linked to innovations in technology that may affect the operation and growth at the industry and national economic level. This includes technology incentives, the level or degree of innovation, automation opportunities, and new advancement in research and development (R&D) activity. Changes to existing technology and communication infrastructures must also be considered.

Political Factors are governing influences generated by changes in the country's national and local government policies towards corporate trade, taxation, and competition (anti-trust) regulations. Factors that determine the stability of the political climate includes the level of government activity and the scope of laws that are enacted. Legal factors are political subcomponents that deal with compliance related issues, such as legal restraints, federal regulations, and the upholding of health and safety laws that relate to general workplace conditions. Major key political factors that influence the labor market include policy changes to employment law, minimum wage, taxation, and immigration law.

Sociocultural Factors are social influences that encompass the cultural and demographic aspects of the external market environment. This includes population trends and demographic changes that impact the labor supply and demand trends, such as the population growth rate, age distribution, and income distribution. Other Sociocultural factors consider the beliefs, customs, and standards held by different business and society cultures. These factors are especially important for learning more about employer needs/preferences in addition to the local labor force and their willingness to work under certain conditions. Other workforce-related sociocultural factors relate to conditions that impact work-life balance, such as career attitudes, workplace flexibility and changes to employment and healthcare benefits.

3.2 Micro-Environmental Factors

Micro-Environmental Factors are those that are distinct and individual, such as customers, producers and competitors. These factors can be best described by applying Michael Porter's Five Force comparative model to the labor market industry.

Porter's Five Force comparative analysis helps organizations or individuals assess the competition within a given industry [15]. The model is based on the insight that a successful business strategy should meet the opportunities and threats in an organization's external environment. Each of the five forces; competitive rivalry, threat of new entrants, threat of substitutes, buyer power, and supplier power help organizations determine how the intensity of competition will influence an industry's potential for profitability.

Competitive Rivalry represents the level of competition that exists within the employment industry or otherwise thought as the rivalry among qualified candidates. The level of competitiveness depends on the number of candidates and their ability to differentiate themselves from the competition based on experience, skills, or the strength of their network. Factors influencing competitive rivalry include the number of competitors, diversity of skillset, educational requirements, industry labor force concentration, employment mobility, and changes to labor supply.

Threat of new entrants represents new workers entering the job market. The threat may be a factor of population dynamics, changes to immigration policies, or changes in the barrier levels required for entry. Examples of Barrier levels required for entry include educational requirements set by hiring companies, new skill requirements resulting from new innovative technology advances, and legal requirements specific to a particular industry, such as licenses, certifications and/or regulatory requirements. Factors influencing the threat of new entrants include shifts in population demographics adjusting the working age, changes in turnover volume of new entrants, and changes to immigration policy.

Threat of Substitutes represents an organization's ability to substitute an alternative method of labor that can perform the same functions. This may be a result of new disruptive technologies or situations that introduce new potential for automation. It could also be the result of the attractiveness to outsource the role at reduced cost. Factors influencing the threat of new substitutes include recent advances in productivity technology, changes in demand for part-time employee/contractors, changes in part-time and contractor labor cost relative to full-time employee cost, and changes to the supply/demand for outsourcing work offshore.

Buyer Power of Employers consists of organizations looking to increase in size by hiring more employees. "Buyer Power" refers to the balance of power in the relationship between the workers and place of employment. If the worker has unique skills, they will have the power to negotiate higher wages. At the same time, if there are many workers looking for employment and only a few organizations are hiring, those firms will be able to justify offering low wages given the abundance of supply available to them. Factors influencing the buyer power of employers include the current level of education and skills required by organizations that are hiring, changes in labor demand, industry concentration, average organization size, price-wage sensitivity, and employee to employer confidence ratio.

Supplier Power in relation to the labor market, represents the organizations who control the supply of job/career postings through a platform that uses algorithms to connect workers with hiring organizations. These online platform providers have the bargaining power of acting as job/career gatekeepers. Factors influencing supplier power include the size, concentration, degree of importance and the availability of substitutes for career/job related platforms.

3.3 Internal Environmental Factors

Internal environment factors are comprised of internal elements unique to the organization in terms of the benefits offered to the employee workforce. Contrary to the behavior shown in macro-environmental factors, organizations can exercise control and influence possible outcomes. Key internal environment factors include an organization's size, culture, executive leadership, overall organization health, and shareholder stability [3].

4 Data Collection Methodologies

4.1 HR Information Systems (HRIS)

Human Resources (HR) departments of global companies require global databases that capture, process and track workforce information such as employee's attrition and hiring, compensation and benefits, ethnicity, gender, cultural, and

nationality distributions. The data from these On-Line Transactional Processing (OLTP) systems are further distilled and loaded into data lakes and data warehouses for Decision Support and On-Line Analytical Processing (OLAP). Leveraging and applying advanced analytical techniques against the data in these workforce systems provides HR professionals and business leaders with intelligent business insights, ability to predict changes and make informed strategic and tactical decisions [13]. HR departments of global companies assemble data such as demographic, hiring, terminations, talent, compensation, benefits, ethnicity, gender, cultural, and nationality distributions and load the same into data warehouses and data marts for analytical processing.

By analyzing the past and current data, business analysts distill business insights and make fact-based decisions supported by data. The global HR information systems consist of several component systems that are interdependent. The various components may be broadly classified into the following main sub-systems: data warehousing, data analytics, data mining, data mashups and information delivery system. These tools and processes are critical to formulate hypotheses to design data and analytical models to compute and communicate results to appropriate users. These users will then draw business insight from the results and shape business decisions that ultimately will improve organizational performance [13].

The data used in our models has been extracted from the HR system of Company “X”. It has been anonymized for security purposes but has all the intrinsic characteristics required for predicting attrition. The data is collected about employees at hire, updated through their employment lifecycle and finally updated and frozen at their termination from the organization.

4.2 Engagement Surveys

The main aim of conducting an employee engagement survey is to find out the factors that drive employees to perform their best and the ones that can put them off. It is an important tool to communicate and sync between what top management offers and what employees expect or vice versa. Engagement surveys are important because they provide employees a channel for sharing open feedback. It is a valuable channel to establish two-way communication between employees and their managers and senior leadership. Employees feel engaged in the development process providing them a direct voice to the leadership team. According to an HBR article by Scott Judd, Eric O’Rourke and Adam Grant, “Surveys are still great predictors of behavior and are a vehicle for changing behavior. They give employees the chance to be heard and feel that their input is valuable to the organization’s growth”. They further suggest that “People who don’t fill out either of our two annual surveys are 2.6 times more likely to leave in the next six months” [9]. Smart technology and big data further help us figure out what matters most to our people by analyzing the results of the engagement survey. This makes surveys even more important, not less. Companies are investing and leveraging innovative machine-learning algorithms and techniques that crunch big data to measure employee engagement. Some organizations also

resort to email response times and network connections outside one's core team to forecast turnover risk by tracking signals such as how often employees update their resumes. Our data from Company "X" contains employee ratings of their direct managers which is critical in determining the employee mindset vis-à-vis their manager. In addition, they also respond to numerous more direct questions regarding the organization, senior leadership, facilities and if they expect to be at the organization within the next 1,3 or 5 years. The employee responses are protected to ensure the anonymity of the employee and encourage open feedback. Our analysis does not include these responses for reasons of strict confidentiality. However, they can be a very valuable source of information and important contributing factors used for the prediction of attrition.

4.3 Exit Interviews

Exit interviews are another tool in the HR toolbox to better understand why employees chose to leave, though by this point, it is usually too late to retain them. "The purpose of an exit interview is to assess the overall employee experience within the organization and identify opportunities to improve retention and engagement"³. The feedback gained can be leveraged to improve aspects of the organization, better retain employees, and reduce attrition especially regrettable attrition. Establishing a clear set of standards and protocols when conducting exit interviews play an essential role in risk management and eliminating bias – conscious and unconscious. The following story is a case in point on the value of exit interviews - "An international financial services company hired a midlevel manager to oversee a department of 17 employees. A year later only eight remained: Four had resigned and five had transferred. To understand what led to the exodus, an executive looked at the exit interviews of the four employees who had resigned and discovered that they had all told the same story: The manager lacked critical leadership skills, such as showing appreciation, engendering commitment, and communicating vision and strategy. More important, the interviews suggested a deeper, systemic problem: The organization was promoting managers based on technical rather than managerial skill. The executive committee adjusted the company's promotion process accordingly" [17]. An organization's assets in today's knowledge economy are its employees especially its highly skilled employees. Understanding and learning what makes employees tick is essential to the success of the organization. Therefore, companies must dedicate time and resources to learn - why they stay, why they leave, and how the organization is perceived, and the changes needed. A thoughtful exit interview process can create a constant flow of invaluable feedback. Many large organizations use an external service provider for this service to further permit the departing employee to be open and transparent. Company X however conducts these exit interviews in-house. Unfortunately, at the time of conducting our analysis, this data was not contained within the HR information system.

³ Additional information can be found at <https://www.hracity.com/blog/importance-of-exit-interviews>. Last accessed 24 Jul, 2020.

5 Model approaches and evaluation

Supervised machine learning techniques are now widely used in business applications across industry practices to predict outcomes. In our research, we explore some of these methods help predict employee turnover. We will explore the factors that contribute toward attrition and how to better control the outcome.

5.1 Model Approach

Supervised machine learning is defined as: given a set of objects each labeled with one of k distinct class labels, learn a function f which classifies new objects into one of k classes [16]. The labeled data set used to learn is called the training data. The new data set that gets classified into the k classes is known as test data. The method to find the function f is referred to as the learning algorithm. The function f is the predictive model.

Logistic Regression: Logistic Regression, also known as Logit Regression, is commonly used to estimate the probability of an instance belonging to a particular class [6]. This model computes a weighed sum of input features and a bias. However, instead of providing the label directly, it outputs the logistic of the result. It is useful when the dependent variable is categorical [20]. The form of the model is:

$$P(Y|X, W) = \frac{1}{1 + e^{-(w_0 + \sum_{i=1}^N w_i X_i)}} \quad (1)$$

In order to prevent overfitting, the logistic regression is often used with regularization techniques like Lasso (L1) or Ridge (L2).

Support Vector Machine: Support Vector Machine (SVM) can perform linear or non-linear classification, regression, and outlier detection [6]. SVMs perform well in a variety of settings and are often considered one of the best “out of the box” classifiers [8]. SVM is generally used for classification of complex small or medium sized datasets. The underlying concept of SVM is to define a hyperplane that separates the p -dimensional data into two classes, by choosing the maximal margin hyperplane [8], which is the hyperplane that maximizes the geometric distance to the nearest data points [21].

Naïve Bayes: Naïve Bayes classifiers are very similar to the linear models. It is based on Bayes Theorem that describes the occurrence probability of an event based on prior knowledge of related features. The efficiency comes from learning parameters by looking at each feature individually and collect simple per-class statistics from each feature. It assumes that the presence of a feature would not influence any other features. The derivation function⁴ to obtain the class, given

⁴ Additional details on derivation of Equation 2 can be found on <https://towardsdatascience.com/naive-bayes-classifier-81d512f50a7c>. Last access 25 Mar, 2020.

the predictors is:

$$y = \operatorname{argmax}_y P(y) \prod_{i=1}^n P(x_i|y) \quad (2)$$

K-Nearest Neighbors: The KNN algorithm works by identifying K data points in the training data that are closest to the new instance and classifies it based on the majority vote of these neighbors or is inversely proportional to the distance computed. The KNN classifier estimates [8] the conditional probability for class j as the fraction of point in N_0 whose response values equal to j:

$$P(Y = j|X = x_0) = \frac{1}{K} \sum_{i \in N_0} I(y_i = j) \quad (3)$$

N_0 represents the K closest points in the training data closest to the new instance x_0 . While building a tree the Gini index or the cross-entropy are used to evaluate the quality of a split as they are sensitive to node purity.

While decision trees are easy to explain and mirror human decision making, they are not very robust. A small change in the data can cause a large change in the final estimated tree. Also, trees don't have the same level of predictive accuracy as some of the other approaches [8].

Decision Tree and Random Forests: Decision trees can be applied to both regression and classification problems. While a regression tree is used to predict a quantitative response, a classification tree is used to predict a qualitative one. This algorithm constructs a tree from a training dataset in which each node is an attribute and branches are the corresponding values

Decision trees form the basis for Random Forests. Random Forest is an ensemble of decision trees trained via the bagging method. The Random Forest algorithm introduces randomness when trees grow and rather than search for the best feature when splitting a node, it searches for the best feature among a random subset of features [6]. It yields a better model by trading a higher bias for a lower variance. Random Forests are also very helpful to get a good understanding of feature importance when we need to perform feature selection.

Ensemble Models: Ensemble models combine the outcome of multiple models to improve the overall performance of the classification. It is the process of executing two or more models and synthesizing the results so that the model output is based on the classification accuracy of more than one model.

In our analysis, the following ensemble models were executed with different combinations of meta classifiers:

1. Vecstack
2. StackingCVClassifier
3. StackingClassifier
4. VotingClassifier

5.2 Data pre-processing and model evaluation

Model evaluation helps us identify our best model and benchmark how the chosen model would work in the future. Using the below techniques would help us guard against overfitting our data. Overfitting is a phenomenon in which our model would work exceptionally well for the test data but would fail to predict reliably for unseen data. The techniques discussed below will help generalize the model to changing data sets.

Train-Test: This technique provides us a way of voluntarily holding back part of the data to test whether the model works. If we use our entire dataset to train the model, then the model will always predict the correct category for any entry in the data set [14]. This process of setting aside a part of the dataset voluntarily to evaluate the model stops the model from being too optimistic when predicting the outcome. Training the model on the entire data set could also lead to data snooping bias [6]. This kind of bias results from refining too many parameters to improve the model's performance on a data set. In our analysis, we will split the data into training and test data sets (80:20) randomly.

Cross-Validation: This method is a statistical method of evaluating generalization performance that is more stable and thorough than using a split into a training and test set [14]. In cross-validation, the data set is split multiple times and the model is trained on each of these splits. For our analysis, we will be using stratified k-fold cross-validation. The folds are selected so that the mean response value is approximately equal in all the folds. In case of a dichotomous classification, each fold will contain roughly the same proportion of the two class labels.

Balancing the dataset (Upsampling): The attrition dataset is asymmetric and the model accuracy tends to be biased towards to the majority class. So, we synthetically balanced the dataset using the SMOTE method to upsample the minority class "Status = 1". Using this method, the number of observations in the minority class was increased to match the majority class. There was no change to the observations in the majority class. Balancing the dataset lets us use the Accuracy score along with Recall and Precision to compare model performance.

Feature Scaling: As part of preprocessing, we reduce disparate feature scales by normalizing the values. This helps machine learning classifiers perform better so that a feature with greater scale is not treated as more important by the model. For example, in HR datasets, a feature like Hire Year would be between 1985-2020 whereas Tenure may be categorized into 5-10 year buckets.

Feature Importance: The analysis explored the features that had a greater influence on the attrition within the organization. Understanding the features' influence on attrition helps organizations curb the attrition rate and retain high value employees. We used the Random Forest classifier to list the top most 10 important features along with their feature importance score.

The feature importance plot (see Fig. 1) is generated before any preprocessing steps like Scaling or One-hot encoding. The most significant features identified were, 1) Talent Assessment rating; 2) Hire Year; 3) Allocation; 4) Employee Category (Manager/Employee); 5) Manager ID (indicating that there is correlation between attrition and Manager); and 6) Job Level.

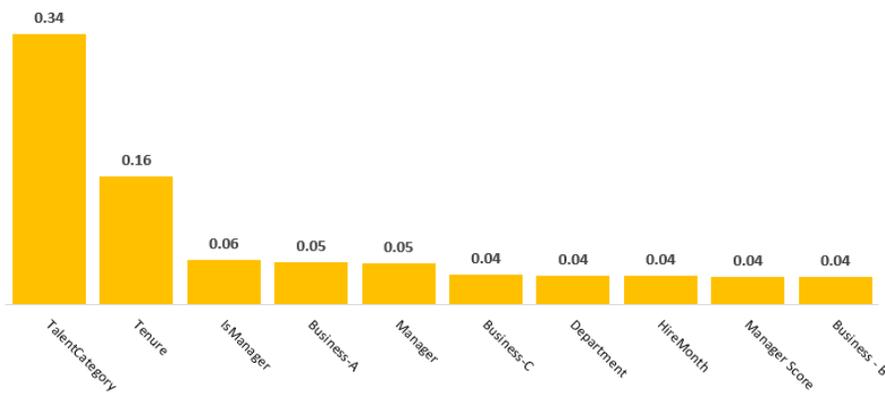


Fig. 1. Top 10 Significant Features and their Importance Scores

6 Model Effectiveness

In our dataset, the proportion of people staying in the organization are far greater than those who leave. This would cause an imbalanced dataset. In such scenarios, just looking at the accuracy is not a reliable option [21]. So, along with Accuracy, we also look at Recall, F1-Score, and area under the Receiver Operating Characteristic curve (ROC-AUC). We'll provide a brief overview of each of the model effectiveness measures. Refer to Table 1 to understand the definitions provided below.

Accuracy: Accuracy is the ratio of correctly classified predictions to the total observations. While accuracy is a great measure to evaluate the model efficiency, it is reliable only when we have symmetric data sets. Therefore, we must look

Table 1. Accuracy, Recall, and F1-Score

	Predicted Class=Yes	Predicted Class=No
Actual Class=Yes	True Positive (TP)	False Negative (FN)
Actual Class=No	False Positive (FP)	True Negative (TN)

beyond just the accuracy to evaluate our model performance.

$$Accuracy = (TP + TN)/(TP + FP + TN + FN) \quad (4)$$

Precision: Precision is the ratio of correctly predicted positive predictions to all the positive predictions made by the model. The question precision answers is: of all the employees labeled as attrition, how many actually left? Precision is appropriate when minimizing false positives in our model. Even though we are not using this metric for evaluation, it is important to understand this to calculate the F1-Score.

$$Precision = TP/(TP + FP) \quad (5)$$

Recall: Recall also known as sensitivity, is the ratio of correct positive predictions to all the observations in actual class=Yes. The question recall answers is: of all the attrition, how many did we label correctly? Recall is appropriate when minimizing false negatives in our model.

$$Recall = TP/(TP + FN) \quad (6)$$

F1-Score: Both precision and recall provide us only half the information. F1-Score is the weighted average of precision and recall. In a way, this measure provides us a method to capture precision and recall into a single measure that captures both properties. In simpler terms, it takes both false positives and false negatives into consideration while calculating the model performance. This is the most common metric for imbalanced classification problems.

$$F1Score = 2 * Recall * Precision / (Recall + Precision) \quad (7)$$

ROC-AUC: The Receiver Operating Characteristic (ROC) curve is common model evaluation tool for binary classifiers. It is similar to the concept of plotting the precision vs. recall. However, the ROC curve plots the true positive rate (another name for recall) against the false positive rate (FPR) [6]. The FPR is the ratio of negative instances that are incorrectly labeled as positive. The model performance is determined by the area under the curve. Better models have more

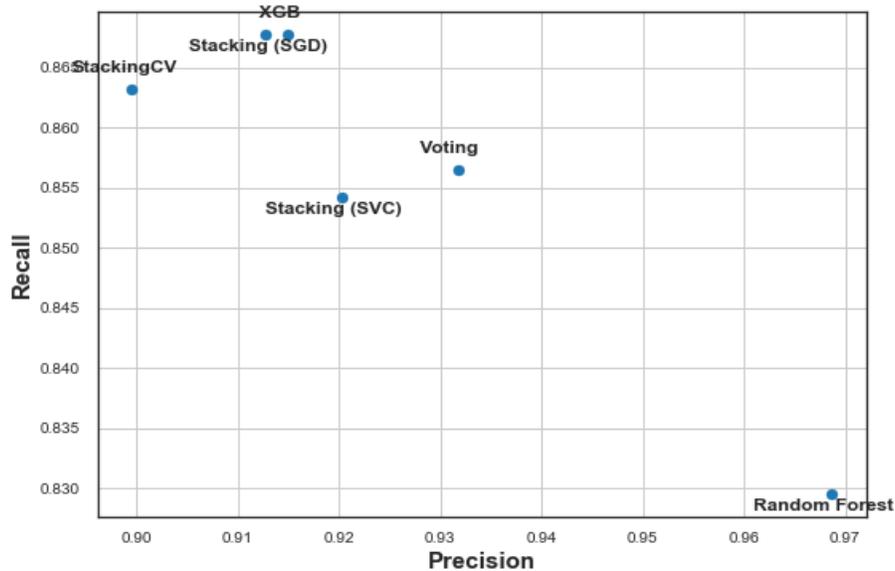


Fig. 2. Precision vs. Recall Scatterplot

area under the curve. ROC-AUC is one of the most important evaluation metrics for checking any classification model's performance. Higher the AUC, better the model is at distinguishing between employees likely to leave vs. stay with the organization.

To assess the accuracies of the model, the ROC curves of all the models were plotted (see Fig. 3). Since we executed the model with multiple parameters and features, we sorted the models by AUC scores and selected one output per model based on the AUC score. While the second ensemble model has the best AUC value, other ensemble models are not far behind. The Random Forest and XGBoost models complete the top-5 list with all of them having a score of 0.959 or better. Generally, models with 0.9 or better ROC-AUC score are generally said to be able to provide outstanding classification of the target variable. 9 of our top 10 models had a ROC-AUC score greater than 0.9.

Comparing Model Effectiveness The models were executed with various flavors of parameters and features. Most of the models were executed for a combination of all vs. reduced features and default vs. best parameters. After the execution, only 1 instance of the model was retained in the result set based on the best AUC score.

Table 2 provides the model effectiveness parameters for all the models executed on the dataset and is sorted from the best to worst AUC scores. Since we were trying to reduce False Negatives, we were striving to improve the re-

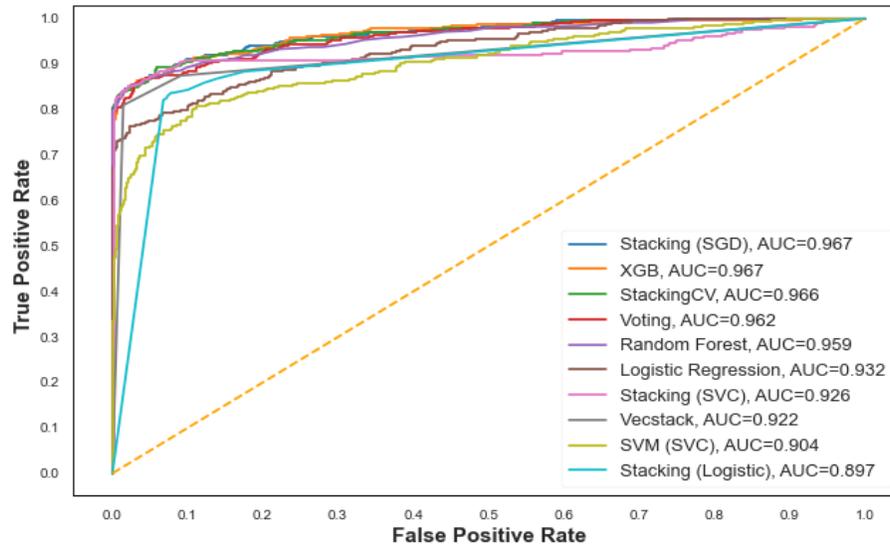


Fig. 3. Best ROC Curves by Classifier

call score. Though the Gaussian Naïve Bayes model had the best recall score, it had the worst accuracy and precisions scores. In general, the Ensemble models performed better than the independent models. The Stacking Ensemble model was marginally better than the XGBoost model in terms of AUC and recall score but was pipped marginally in terms of precision score. The performance of the Random Forest model was comparable to the top 3 ensemble models and the XGBoost model. Logistic Regression had a good AUC Score but both the precision and recall scores were lower than 0.8.

Table 2. Classifier Effectiveness

Classifier	Accuracy	Recall	Precision	AUC Score
Stacking (SGD)	0.929	0.868	0.913	0.967
XGB	0.929	0.868	0.915	0.967
StackingCV	0.923	0.863	0.9	0.966
Voting	0.932	0.857	0.932	0.962
Random Forest	0.935	0.83	0.969	0.959
Logistic Regression	0.866	0.798	0.798	0.932
Stacking (SVC)	0.927	0.854	0.92	0.926
Vecstack	0.898	0.874	0.826	0.922
SVM (SVC)	0.861	0.78	0.796	0.904
Stacking (Logistic)	0.894	0.836	0.842	0.897
Bernoulli Naive-Bayes	0.815	0.791	0.694	0.889
Decision Tree	0.883	0.868	0.798	0.879
Gaussian Naïve-Bayes	0.524	0.982	0.409	0.655

7 Conclusion

The model identified the most important features that have the greatest contribution to the classification of attrition. It is abundantly evident that Talent Assessment which is the stratification of talent is a key differentiator in the stratification of talent in organization “X”. This driver also has a major impact on classification of attrition at the employee level. Further analysis at the organizational level will need to be conducted to determine why certain talent categories have a greater propensity to voluntarily leave the organization than others. The features identified as important by the model will help the HR practitioners in their analysis. In addition to TalentAssessment, HireYear also emerged as an important feature. HireYear is synonymous with the tenure of an employee and is an important driver in an employee’s decision to leave the organization. Deeper patterns regarding tenure will need to be analyzed further to better understand which sub-groups within this factor lean more to leaving the organization than others. Another factor is A-Alloc which represents the percentage time of their work allocated to business A by the employees at this organization. It provides the HR practitioner an indicator that greater attention needs to be focused on individuals with time allocations to this business. The reasons could vary as to why this specific business is causing employees allocated to this business wanting to leave the organization. It necessitates a much deeper analysis within that business. The IsMgr and ManagerID are closely related to the Manager and are also important factors in this model. The IsManager variable indicates if the employee is a manager or an individual contributor. There is value in understanding if the organization is potentially going to lose more of its managers or just individual contributors. Additionally, a deeper analysis into why this is and what can be done to retain either group by delving into the challenges may be of

value. As the saying goes, “Employees don’t leave Companies, they leave Managers”. “A Gallup poll of more 1 million employed U.S. workers concluded that the No. 1 reason people quit their jobs is a bad boss or immediate supervisor. 75% of workers who voluntarily left their jobs did so because of their bosses and not the position itself. Despite how good a job may be, people will quit if the reporting relationship is not healthy. “People leave managers not companies...in the end, turnover is mostly a manager issue.”⁵.

On the flip side of this argument are also other studies that indicate it may not be the managers but the work. While people are more likely to jump ship when they have a horrible boss, at Facebook many employees said they were happy with their managers and their decision to leave was based on their work. They left when their job was not very enjoyable, their strengths were not being used or they were not growing their careers. According to the HBR article “Working with our (Facebook) People Analytics team, we crunched our survey data to predict who would stay or leave in the next six months, and in the process, we learned something interesting about those who eventually stayed. They found their work 31% more enjoyable, used their strengths 33% more often, and expressed 37% more confidence that they were gaining the skills and experiences they need to further develop their careers. This highlights three ways that managers can customize experiences for their people: enable them to do work they enjoy, help them play to their strengths, and carve a path for career development that accommodates personal priorities”⁶.

Other studies and articles exist that contradict the fact that employees leave managers and not companies. According to Didier Elzinga Founder and CEO, Culture Amp⁷ it is the biggest lie in HR that people leave managers not companies. Dr. Jason McPherson, Culture Amp’s Chief Scientist tested the veracity of this claim and his findings can be summarized as follows:

1. Yes, people leave bad managers, but it is not the number one reason why people leave an organization
2. In “good” companies, managers make a difference
3. In “bad” companies, good or bad managers make little to no difference to a person’s decision to leave

With a Precision of 94% (correctly predicted positive predictions) and a Recall of 84% (minimizes false negatives) the developed model would greatly assist the HR department of Company “X” in better understanding who are the employees that are likely to leave the organization. This would enable them to adopt mitigating strategies to retain talent they believe is critical and strategic based on the results of the classifier model. Additionally, the model’s prediction

⁵ More information may be found at <https://www.linkedin.com/pulse/employees-dont-leave-companies-managers-brigitte-hyacinth/>. Last accessed 19 Jun 2020.

⁶ More information may be found at <https://hbr.org/2018/01/why-people-really-quit-their-jobs>. Last accessed 19 Jun 2020.

⁷ More information may be found at <https://www.cultureamp.com/blog/the-biggest-lie-in-hr-people-quit-managers/>. Last accessed 19 Jun 2020.

results would enable Company “X” to develop appropriate talent pipelines well in advance. These actions and strategies help shrink the recruiting time to backfill open positions and also consider “top grading” certain roles to further align to business strategy. These are the benefits of the classification model developed to solve a real business issue.

8 Ethical Considerations

8.1 Subjectivity, Discrimination, Subjectivity, and Bias

The transparency of the predictive modeling process and its respective underlying data is crucial for producing outcomes that are effective, fair, and are free from discriminatory negligence. The very nature of most commonly applied algorithms, if applied to humans, is applying prejudice [5]. This occurs because algorithmic models are based on existing data and thus the model will learn from the data and any historical bias embedded in the data. For example, in 2018, Amazon discovered that its algorithm used for hiring decisions was discriminating against women. The model was created using historical job performance data, when males had been the best performers and majority of applicants. As a result, the algorithm gave white males higher scores than females when the gender variable was present. After removing the gender variable, the model continued to exclude or penalize applicants with attributes associated with female candidates (i.e. women’s studies) [18].

Title VII of the Civil Rights Acts of 1964 prohibits employment discrimination in employment based on two conditions, “demonstrable intent to discriminate” and “disparate impact” [10]. This happens to organizations that do not perform the necessary data validation checks required to reduce the risk of making decisions that may have a disparate impact due to biased or non-representative training data, even if the sensitive variables, such as gender and age, are no longer present. Other implications include lack of transparency and data unambiguity. Organizations can reduce the risk of potential bias and/or discrimination by:

- Proactively test and adjust as needed for any known or unknown historical bias that may be embedded in the data.
- Continuously monitor for new unforeseen bias as part of the standard modeling process (i.e. enable model to make dynamic adjustments)
- Understand which variables are comparable and which are unique to certain constraints, such as geographical location, primary legal jurisdiction, etc. For example, the rate of US compensation will be unique to the US and cannot be compared to others without taking into account all possible confounding variables.

8.2 Data Privacy Abuse

Data controllers and processors of employee personal data are required legally by most countries to place appropriate technical and organizational measures to

implement data protection principles. Business activities that handle personal data must be designed to and built with consideration of the General Data Protection Regulation (GDPR) principles and provide safeguards to protect data. Failure to comply with global regulations increases the risk of data abuse. Best practices for reducing risk of data abuse include:

- Dis-identify personal employee data by removing the 18 HIPAA defined personal identifiers or fully anonymize data where appropriate.
- Collect and maintain only necessary data.
- Utilize data only for the purposes that it was intended for.
- Create a data governance structure that includes all primary stakeholders including an employee representative (i.e. employee union rep, employee representative, and/or independent risk auditor) responsible for supervising the construction, deployment, and monitoring of predictive models.

8.3 Negligent Corporate Duty of Care

The impersonal nature of predictive analytics can sometime lead to business organizations prioritizing company health over the best interest of the individual employee. For example, it is well documented that one of the leading software vendors for HR, generates machine learning algorithms that score individual employees based on changes to their Linked In profiles and/or other job-related social media posts. To reduce corporate abuse many countries have established core policies that promote basic ethical and legal standards for business organizations operating within legal jurisdiction. US Public Corporations along with EU Corporations that fall under the GDPR must exercise undue care to the health and wellbeing of their employees. Predictive results from the automated processing of sensitive personal data may only generate positive legal effects or similar significant outcome that benefits the data owner. Organizations must have a clearly defined purpose for utilizing employee personal data that does not pose harm to the data owner.

Best practices for preventing undue harm to employees:

- Ensure that primary outcomes are measurable and center around the employee
- Do not use model outcomes or data monitoring to furtively inform performance reviews.
- Do not use any personal data without receiving consent from the employee providing the data.

References

1. Abelson, M.A.: Strategic management of turnover: a model for the health service administrator. *Health Care Management Review* **11**(2), 61–71 (1986)
2. Aguilar, F.J.: *Scanning the business environment*. Macmillan (1967)

3. Bright, D.S., Cortes, A.H., Hartmann, E., Parboteeah, K.P., Pierce, J.L., Reece, M., Shah, A., Terjesen, S., Weiss, J., White, M.A., et al.: Principles of Management. OpenStax, Rice University (2019)
4. Fernández-Aráoz, C., Groysberg, B., Nohria, N.: How to hang on to your high potentials (2011)
5. Garzcarek, U., Steuer, D.: Approaching ethical guidelines for data scientists. In: Applications in Statistical Computing, pp. 151–169. Springer (2019)
6. Géron, A.: Hands-On Machine Learning with Scikit-Learn, Keras, and Tensor-Flow: Concepts, Tools, and Techniques to Build Intelligent Systems. O'Reilly Media (2019)
7. Goswami, B.K., Jha, S.: Attrition issues and retention challenges of employees. International Journal of Scientific & Engineering Research **3**(4), 1–6 (2012)
8. James, G., Witten, D., Hastie, T., Tibshirani, R.: An introduction to statistical learning, vol. 112. Springer (2013)
9. Judd, S., O'Rourke, E., Grant, A.: Employee surveys are still one of the best ways to measure engagement. Harvard Business Review (2018)
10. Koffie-Lart, D., Tyson, C.J.: Title vii of the civil rights act of 1964. Geo. J. Gender & L. **6**, 615 (2005)
11. Marsden, T.: What is the true cost of attrition? Strategic HR Review (2016)
12. Martin, J., Schmidt, C.: How to keep your top talent. Harvard business review **88**(5), 54–61 (2010)
13. Mora, G., Kapoor, B., Sherif, J.: Global human resources (hr) information systems. Kybernetes (2012)
14. Muller, A.C., Guido, S.: Introduction to machine learning with Python: a guide for data scientists. O'Reilly Media (2017)
15. Porter, M.E.: The five competitive forces that shape strategy. Harvard business review **86**(1), 25–40 (2008)
16. Saradhi, V.V., Palshikar, G.K.: Employee churn prediction. Expert Systems with Applications **38**(3), 1999–2006 (2011)
17. Spain, E., Groysberg, B.: Making exit interviews count. Harvard Business Review **94**(4), 88–95 (2016)
18. Tambe, P., Cappelli, P., Yakubovich, V.: Artificial intelligence in human resources management: Challenges and a path forward. California Management Review **61**(4), 15–42 (2019)
19. Yadav, S., Jain, A., Singh, D.: Early prediction of employee attrition using data mining techniques. In: 2018 IEEE 8th International Advance Computing Conference (IACC). pp. 349–354. IEEE (2018)
20. Yedida, R., Reddy, R., Vahi, R., Jana, R., GV, A., Kulkarni, D.: Employee attrition prediction. arXiv preprint arXiv:1806.10480 (2018)
21. Zhao, Y., Hryniewicki, M.K., Cheng, F., Fu, B., Zhu, X.: Employee turnover prediction with machine learning: A reliable approach. In: Proceedings of SAI intelligent systems conference. pp. 737–758. Springer (2018)