
HR Attrition Analysis

Anja Wu

Executive Summary

Using an employee dataset from Kaggle, factors affecting attrition were studied, along with creating a model to predict attrition based on the employee features contained within the dataset. Several models and methods were tried to get maximum accuracy and deal with the imbalance of attrition feature. The best model was the decision tree, primarily when the oversampling method was used to randomly distribute the training data. A logistic regression was used to further understand any relationships between features and attrition. The top features which had the largest effects on attrition were: total working years, years at company, years with current manager, and percent salary hike.

1 Introduction

I am fairly new to analysis and have not worked with any HR analysis prior to this project but I will be looking at attrition in a company with specific employee features. Through the analysis of employee data using a Kaggle dataset [1], I will be trying to predict attrition of employees within a year at a company based on the data given. This dataset drew my interest given articles I have read online about the Great Resignation phenomenon that has arisen during the pandemic. The main goal is to understand what factors influence attrition the most and if there are any suggestions for the company to try to keep more of their employees. A secondary goal will be to find the best predictive model.

Two part hypothesis:

1. There are several factors which will impact attrition the most: "EnvironmentSatisfaction", "JobSatisfaction", "PercentSalaryHike", and "YearsWithCurrManager". I think the best model to give me the relationship between the features and attrition is logistic regression as it allows the analysis of each individual features contribution to the outcome.
2. I think the decision tree model will lead to the best predictive model because of research that indicated it performs particularly well with imbalance outcomes.

2 Problem Definition

Problem: Attrition in general is a difficult topic to predict because so many factors need to be considered. Even if you have the perfect model with all factors considered, people's intent to leave changes on a variety of things that is based on personal preferences. This can make it difficult to extract a general rule for all people. Due to the vast variety of data in social science (such as this type of data), you would not expect the same level of correlation that you would want to see in other situations. It is enough to have a moderate correlation between any of your features and attrition to draw conclusions. It is all relative to the other features you have and how they perform. Analyzing attrition by looking at the correlations can help businesses adjust their company structure to retain as many personnel as possible because it gives more insight into the why.

Dataset: The dataset has personal information about the employee (age, gender, marital status, distance from office, etc.), work background of employee (education, number of companies worked, etc.), employee job information (job level, role, job performance, etc.), and job pay (monthly income, stock options, etc.). The dataset also had whether the employee has left in the past year or not (attrition). When looking at attrition, not all attrition is the same; it is important to look at regrettable

vs non-regrettable attrition. Because sometimes attrition can be good, if it is for people who aren't in the right role to make the impact on what they are working on. This is something that should be looked at. This particular dataset only had job performance of 3 or 4, which is fairly high. So I went ahead with the assumption that they were only looking at their higher performers and what was causing them to leave.

Data Pre-Processing: During the pre-processing of data, there were several things found:

- Several columns had only one possible option, this means they do not add any signal towards the analysis of attrition. I deleted these columns, so as to reduce dimensions.
- One column (BusinessTravel) had ordinal categories and was encoded ranging from 0 to n-1, in order. The other columns that had features that were categorical were encoded using one hot encoding.
- When doing regressions, collinearity must be checked. The correlation matrix can be seen to have several variables that have a high correlation. The best way to deal with this is to check for multicollinearity using VIF. [5] [3]
 - VIF runs a linear correlation for each feature against all other features
 - For the VIF calculations it can be seen that all values are below 5 implying that features are not highly correlated. This means that we can proceed with our regression without dropping any columns.

Dealing with imbalance: Upon further exploration of the data, it was discovered that the attrition feature is imbalance (which makes sense). There are different methods to deal with imbalance and I found a good article [2] that summarizes 10 ways. I chose the following methods:

- a) Random over-sampling with imblearn: The imblearn library takes a random sample from the minority class, with replacement, until the number of samples is the same as the majority class.
- b) Random under-sampling with imblearn: The imblearn library takes a randomly generated sample of the majority class, in my case, without replacement, to match the number of samples in the minority class.
- c) Synthetic Minority Oversampling Technique (SMOTE): In general, SMOTE is an oversampling technique which uses k-nearest neighbours to increase the sample of minority class. It does this by taking k-nearest neighbours from a randomly selected minority point and then randomly selecting one of the neighbours. Then it creates a new synthetic point in between the minority point and its nearest neighbour. For this, I chose $k=5$. I tried a bunch of different k-values multiple times (5, 10, 15, 25, 30, 50, 100) and they all gave relatively the same predictive power. So for computational power and not risking overfitting, I chose the lowest k-value.
- d) Under-sampling: Tomek links: In general, Tomek links uses k-Nearest Neighbours to select for undersampling. The concept is that the algorithm will pair minority and majority points on the boundary that are a small distance from each other and will convert the majority label to the minority label. As can be seen in our data, the number of points that it converted was very minor, leading to still a great divide between the majority and the minority label. So this method was not a great one to be used for this dataset.

Picking a model: Through the course we learnt k-nearest neighbours, decision trees, random forests, linear regression, logistic regression, and forward-feed neural networks. In my initial goal of this project, seeing which features have the greatest effect on attrition, a regression would be the best bet. Regressions allow the analysis of the relationship between variables. A logistic regression is used when your outcome is binary and you would like to know to what degree your features affect the outcome. I wanted to approach this problem to maximize the understanding of what factors lead to attrition, so I chose the logistic regression algorithm. Regression models allow the understanding of which features contribute the most to the outcomes you are looking for, this is done by looking at the coefficients. The accuracy of predictions was going to indicate to me how well the model predicted the importance of the features. A secondary goal of this project was to get as good of a predictive model to be able to predict whether an individual would leave the company based on the features inputted.

Through research, I found that decision trees, in particular, usually perform the best with imbalance data. However, I still wanted to look at the relationship between the features and attrition, so I went

ahead and did a logistic regression first. Then I also wanted to look at the best approach to see the difference in the performance, so I also did a decision tree. Random forest ensemble method would be a good idea to also use, however, I had found several drawbacks during my research[4] [6] including being more of a black box, and having very little interpretability. So I discounted this method, due to the fact that a decision tree was said to perform quite well for imbalanced data and is slightly more interpretable than random forest.

3 Model

Below I will be discussing how I designed the two models, keeping in mind that for all I used 5 different sampling methods to deal with the imbalanced data: no adaptation (this creates the baseline), random over-sampling with imblearn, random under-sampling with imblearn, Synthetic Minority Oversampling Technique (SMOTE), and under-sampling using Tomek links.

3.1 Model 1: Logistic Regression

Due to the correlation that I had found during the data exploration phase, I chose to try different variations of dropping correlated columns to analyze effects on the attrition outcome. Overall, I found not much of a difference when columns were dropped. In fact, the full dataset performed better at predicting than having any potentially correlated features removed. So for all other methods of dealing with imbalance, I used the full dataset.

The data was separated into just training and testing sets, and the distribution of the labels was checked. After I fit the model using scikit LogisticRegression() on the training set, I plotted the correlation coefficients for all features to get more of a visual of what features seemed to have the greatest effect. A confusion matrix was created to display the precision, recall, F1-score, and accuracy of the predictions. After this was done on all 5 methods (listed above), I created a table which held all results (highest correlation coefficient features and the metrics associated with the methods) to be able to more easily compare the different approaches. I also analyzed the most accurate model using statsmodels logistic regression as it provides a p-value and would allow me more understanding of the data.

3.2 Model 2: Decision Tree

I modified the split data function we had in assignment two to work for this dataset. The function splits the data into 3 sets: training, validation, and test. The hyperparameters will be tuned using the training and the validation set, to only be tested once on the test set once the hyperparameters have been chosen. A select model function has been created (similar to assignment 2) which fits the tree onto the training data for various depths, and criterion and returns the accuracy (using the get_acc function) for both the training and validation set.

Like before, the distribution of the attrition labels was analyzed for each method. Then the select model function was used to run different depths, using both Entropy and Gini criterion. The validation set accuracy was plotted along with the training set accuracy for those sets of depths and criterion. This was run 10 times to find an approximate best depth and criterion for each sampling method to be considered the best model (in the code, after each method, there is a description of why each was chosen). I fit the model using DecisionTreeClassifier() on the training set. Like before, a confusion matrix was created to check the accuracy, precision, recall, and F1-score for the final model chosen. Each method had a decision tree created.

4 Results and Findings

4.1 Model 1: Logistic regression

As can be seen in table 1, the method that lent the best predictions was the random oversampling method (Method 1) because it has the best metrics (fair accuracy and has the best precision and recall).

Table 1: Looking at the precision, recall, F1-score, and accuracy of "yes" attrition and "no" attrition predictions completed for each of the sampling methods discussed (Method 0: control, Method 1: Random Oversampling, Method 2: Random Undersampling, Method 3: SMOTE, and Method 4: Tomek Links) for the logistic regression.

	Method0	Method1	Method2	Method3	Method4
no_precision	0.827664	0.643617	0.545455	0.602273	0.816327
yes_precision	0.000000	0.653846	0.000000	0.585052	0.000000
no_recall	1.000000	0.657609	1.000000	0.568365	1.000000
yes_recall	0.000000	0.639785	0.000000	0.618529	0.000000
no_f1	0.905707	0.650538	0.705882	0.584828	0.898876
yes_f1	0.000000	0.646739	0.000000	0.601325	0.000000
accuracy	0.827664	0.648649	0.545455	0.593243	0.816327

Table 2: Statistically significant ($p\text{-value} < 0.05$) logistic coefficients in order from largest absolute value to smallest for the Random Oversampling method (Method 1).

	index	coef	pval
0	is_hr	1.342707	1.079403e-22
1	PerformanceRating	0.892007	5.338988e-19
2	is_single	0.874228	6.040794e-43
3	BusinessTravel_encoded	0.864598	1.706523e-56
4	is_manufactorDir	-0.765492	1.783974e-11
5	is_manager	-0.484653	2.496368e-04
6	is_researchDir	0.347567	6.483611e-03
7	JobSatisfaction	-0.314809	5.501484e-33
8	EnvironmentSatisfaction	-0.302809	6.637195e-33
9	is_healthRep	-0.232032	4.319277e-02
10	WorkLifeBalance	-0.223176	8.203889e-10
11	is_divorced	-0.216213	6.661937e-03
12	YearsWithCurrManager	-0.192712	2.505510e-42
13	is_rd	0.179868	4.411294e-03
14	YearsSinceLastPromotion	0.140197	5.988923e-29
15	TrainingTimesLastYear	-0.140027	1.268333e-09
16	NumCompaniesWorked	0.138259	3.663487e-28
17	TotalWorkingYears	-0.071009	2.940604e-23
18	YearsAtCompany	0.048906	3.225460e-06
19	PercentSalaryHike	-0.025886	2.987233e-02
20	Age	-0.022643	2.932849e-08

Table 2 shows the ranking of the statistically significant logistic regression coefficients. As can be seen:

- is_hr: a person belonging to the HR department is more likely to leave when compared with a person in the sales department.
- PerformanceRating: a person with a higher performance is more likely to leave.
- is_single: a person who is single is more likely to leave when compared with a person who is married.
- BusinessTravel_encoded: a person who travels more is more likely to leave.
- is_manufactorDir: a person who is a Manufacturing Director is less likely to leave than a person who is a Sales Executive.
- is_manager: a person who is a Manager is less likely to leave than a person who is a Sales Executive.
- is_researchDir: a person who is a Research Director is more likely to leave than a person who is a Sales Executive.
- JobSatisfaction: a person who is less satisfied with their job is more likely to leave.
- EnvironmentSatisfaction: a person who is less satisfied with the environment is more likely to leave.
- is_healthRep: a person who is a Healthcare Representative is less likely to leave than a person who is a Sales Executive.
- WorkLifeBalance: a person who has a worse work-life balance is more likely to leave.

- `is_divorced`: a person who is divorced is less likely to leave when compared with a person who is married.
- `YearsWithCurrManager`: the shorter amount of time the person has been with their manager the more likely they are to leave.
- `is_rd`: a person belonging to the R&D department is more likely to leave when compared with a person in the sales department.
- `YearsSinceLastPromotion`: the longer it has been since a person has been promoted the more likely they are to leave.
- `TrainingTimesLastYear`: the more times a person has taken training in the last year, the less likely they are to leave.
- `NumCompaniesWorked`: the more companies a person has worked at the more likely they are to leave
- `TotalWorkingYears`: the less time someone has worked, the more likely they are to leave
- `YearsAtCompany`: the longer someone is at the company, the more likely they are to leave
- `PercentSalaryHike`: the higher the percent salary hike, the less likely a person is to leave.
- `Age`: the older a person is the less likely they are to leave.

This model did a good job in showing the relationship between our features and the attrition outcome. The next model, we will be trying to increase prediction - instead of focusing on the inference.

4.2 Model 2: Decision Tree

As can be seen in table 3, all methods for a decision tree are significantly better at predicting than the logistic regression. In particular, method 1 (random oversampling) performed the best, followed closely by method 3 (SMOTE). It seems that in general the oversampling is what made the predictions better. In analyzing the decision trees created, it was very clear that the data is not sorted according to the most correlated feature to attrition, but rather the algorithm focuses on the most information gain at each split. The algorithm cares most about the final output/classification and not how it gets there, and there could be cases where the optimal tree isn't the one that learns the most at the beginning of the tree (splits at hierarchy don't necessarily signify importance). So this model would not be best at determining relationships between our features and the attrition outcome.

	Method0	Method1	Method2	Method3	Method4
no_precision	0.983636	0.967495	0.915663	0.966102	0.913043
yes_precision	0.875000	0.981293	0.885496	0.967370	0.811475
no_recall	0.974775	0.978723	0.835165	0.971039	0.785047
yes_recall	0.915888	0.971380	0.943089	0.961832	0.925234
no_f1	0.979186	0.973077	0.873563	0.968564	0.844221
yes_f1	0.894977	0.976311	0.913386	0.964593	0.864629
accuracy	0.965257	0.974797	0.897196	0.966697	0.855140

Table 3: Looking at the precision, recall, F1-score, and accuracy of "yes" attrition and "no" attrition predictions completed for each of the sampling methods discussed (Method 0: control, Method 1: Random Oversampling, Method 2: Random Undersampling, Method 3: SMOTE, and Method 4: Tomek Links) for the decision tree model.

Real life application: Finding the relationship between features and attrition can be useful to the company if we are looking at features that they can influence. For things such as total working years and years at the company, not much can be done. However, it would be good for the company to further explore the "why" behind more employees with high salary hikes leaving and if there is anything that can be done to get those individuals to stay (since they are most likely the higher performing employees).

In terms of the predictive capabilities of the decision tree, it might be worthwhile to analyze out of certain employees that are instrumental to your organization, what is the likelihood of them leaving.

Having a model that can predict an individual leaving, would be very useful in knowing where to focus efforts in getting high impact employees who might be close to leaving to stay.

5 Conclusions and Future Work

Overall, logistic regression should be used to find relationships between features and the class (attrition). In this case, although the accuracy of prediction was not high, we can see some relationships between years worked, years spent at the specific company, percent salary hike, and years with the current manager all impact attrition. For the best predictive model, I would recommend using the decision tree with the random oversampling method to deal with imbalance, as it gave the greatest overall accuracy (precision, recall, and F1-score).

Limitations/Future Opportunities:

- The logistic regression coefficients should be transformed using $e^{\text{coefficient}}$ in order to quantify the degree of each feature on the outcome of attrition. This will allow a more in-depth understanding.
- One final limitation is that I do not have much HR experience and I might be missing some key fairness considerations in this domain/context.

Strengths:

- One hot encoding was used, leading to more detailed and accurate analysis of features.
- p-value was observed for the logistic regression allowing a more accurate definition of which variables affect attrition.
- The distributions of the features were checked, which is important for checking for fairness. In general, there seemed to be no major imbalance in the features.
- More than one model and method was used to see which would perform best with the dataset.
- More than one run of each method and model was run to find an average (similar to bagging) for hyperparameters.

References

- [1] V. Choudhary. HR Analytics Case Study, 2018. URL <https://www.kaggle.com/datasets/vjchoudhary7/hr-analytics-case-study>.
- [2] B. Kumar. 10 Techniques to deal with Imbalanced Classes in Machine Learning, 2020. URL <https://www.analyticsvidhya.com/blog/2020/07/10-techniques-to-deal-with-class-imbalance-in-machine-learning/>.
- [3] W.-M. Lee. Statistics in Python — Collinearity and Multicollinearity, 2021. URL <https://towardsdatascience.com/statistics-in-python-collinearity-and-multicollinearity-4cc4dcd82b3f>.
- [4] J. Singh. Random Forest: Pros and Cons, 2020. URL <https://medium.datadriveninvestor.com/random-forest-pros-and-cons-c1c42fb64f04>.
- [5] I. Team. Variance Inflation Factor (VIF), 2021. URL <https://www.investopedia.com/terms/v/variance-inflation-factor.asp>.
- [6] D. Trehan. Why Choose Random Forest and Not Decision Trees, 2020. URL <https://towardsai.net/p/machine-learning/why-choose-random-forest-and-not-decision-trees>.