HIGH SCHOOL EDITION
Journal of Student Research

Election Forecasting Using Macroeconomic and Social Indicators via Machine Learning

Lucas Guan1 and Ganesh Mani[#]

[1] Palo Alto High School, Palo Alto, CA, USA
[#]Advisor

ABSTRACT

A comparative analysis of machine learning models is executed for the forecasting of incumbent party losses during federal elections in democratic countries. A proprietary dataset that encompasses a wide array of potential economic and social factors affecting election outcomes is compiled, and the most significant factors are identified and evaluated. A myriad of the most popular machine learning models for supervised learning are applied to the dataset, utilizing them as classifiers to predict whether the incumbent party stays in power during federal elections for eleven of the world's most populous and democratic countries: the United States, Canada, the United Kingdom, the Netherlands, Austria, Norway, Sweden, Denmark, Australia, India, and New Zealand. The results show that the most significant factors for election outcomes are inflation growth rate, unemployment growth rate, and voter turnout growth rate. Multilayer perceptron produces the most accurate classifications. Additionally, Gaussian models such as Gaussian process classifier and Gaussian naive Bayes have the poorest classification accuracy.

Introduction

Forecasting federal elections has been attempted in long-standing democracies across the world. Although a wide array of research exists on the use of macroeconomic indicators to forecast federal election outcomes, significantly less leverages machine learning. In fact, with recent advances in artificial intelligence and supervised learning, scholarly attention has turned away from economic factors to sentiment analysis using social media platforms. However, the state of a national economy appears to be an important one, as many incumbent parties have lost national elections during times of economic crisis, such as the Canadian Progressive Conservatives in 1993, the Republican Party in the United States in 2008, and the Greek New Democracy Party in 2015. This research applies ten of the most common supervised learning models to predicting democratic election outcomes using macroeconomic and social indicators. To ensure geographical diversity and comparability in the countries studied, 11 of the world's most long-standing, developed, and populous democracies are chosen based on the Democracy Index: the United States, Canada, the United Kingdom, the Netherlands, Austria, Norway, Sweden, Denmark, Australia, India, and New Zealand.

Related Works

Early election predictions focus primarily on the effect of economic conditions on election outcomes. Specifically, the academic debate revolves around conflicting findings on whether economic forces and perceptions play a key role at the ballot box. Lewis-Beck and Stegmaier (2000) find that past research centered on a single country has consistently shown that economic forces exert "a heavy and variegated" impact on democratic elections worldwide. By contrast, Paldam (1991) concludes that economic results "are either insignificant or explain very little" when studying elections between 1948–1985 from a pool of 14 democratic countries. Moreover, Chappell and Veiga (2000), in a study of 13 Western European countries, note that voters only react to increased inflation.

More recent research has shifted from macroeconomic indicators to social sentiments and heavily leverages machine learning. Kennedy, Wojcik, and Lazer (2017) assess the predictability of elections from polling data and achieve 80-90% accuracy across 84 countries. Twitter feeds are another popular source of sentiment data. Tsai et al. (2019) analyze Twitter data to predict the outcomes of the 2018 U.S. midterm elections, Budiharto and Meiliana (2018) for the 2019 Indonesian Presidential election, and Nausheen and Begum (2018) for the 2016 United States Presidential election. This study builds off the previous literature by assessing the significance of a myriad of macroeconomic and social election factors and evaluating the effectiveness of machine learning classifiers for the prediction of democratic election outcomes.

Data

Basic Characteristics

Data on five primary social and economic factors was collected: national economic trends, household purchasing power, voter engagement, social unrest, and global market trends. For each factor, the rate of growth annually and the rate of growth over the entire term of the incumbent, wherever applicable, were calculated. To ensure consistency across countries, all values were expressed as percentages. National economic trend factors consist of one-year gross domestic product (GDP) growth rate, one-year gross national income (GNI) growth rate, incumbent-term GDP growth rate, and incumbent-term GNI growth rate. Household purchasing power factors consist of inflation, one-year GDP per capita growth rate, one-year GNI per capita growth rate, one-year inflation growth rate, incumbent term GDP per capita growth rate, incumbent term GNI per capita growth rate, and incumbent term inflation growth rate. Voter engagement factors are comprised of voter turnout during elections and voter turnout change rate during the incumbent term. Social unrest factors consist of unemployment rate, one-year unemployment growth rate, and incumbent term unemployment growth rate. Global market trend factor is comprised of the MSCI World index annual returns. Table 1 shows a summary of the independent variables used within the research.

Table 1. Input variables collected for each election.

Category	**Variable**
National Economic Trends	One-year GDP growth rate
	One-year GNI growth rate
	Incumbent term GDP growth rate
	Incumbent term GNI growth rate
Household Purchasing Power	Inflation
	One-year GDP per capita growth rate
	One-year GNI per capita growth rate

	One-year inflation growth rate
	Incumbent term GDP per capita growth rate
	Incumbent term GNI per capita growth rate
	Incumbent term inflation growth rate
Voter Engagement	Voter turnout during election
	Incumbent term voter turnout change rate
Social Unrest	Unemployment rate
	One-year unemployment growth rate
	Incumbent term unemployment growth rate
Global Market Trends	MSCI World index annual return

Collection

Data on GDP, GDP per capita, GNI, GNI per capita, inflation, and unemployment is collected from The World Bank. Data from The World Bank is utilized to construct all of the national economic trend, household purchasing power, and social unrest variables. Annual returns of MSCI World index are collected from MSCI. The MSCI World index is chosen because it is the oldest and most comprehensive global stock index, covering companies from all countries studied with the exception of India. Lastly, data on election outcomes and voter turnout is collected from country-specific sources: Elections Canada, United Kingdom Data Service, Electoral Council (Netherlands), Ministry of the Interior (Austria), Statistics Norway, Statistics Sweden, Statistics Denmark, Australian Electoral Commission, Electoral Commission (New Zealand), Election Commission of India, and Gallup (United States).

Variable Construction

While inflation, unemployment rate, and MSCI index annual returns are already processed and ready to use at the time of collection, the remaining 14 variables are manually computed.

Equation 1: We construct the one-year GDP growth rate, GDP_{1y}, where GDP_e is the GDP (in USD) during the election year *e* and GDP_{e-1} is the GDP during the year prior to the election year *e*:

$$GDP_{1y} = ln\left(\frac{GDP_e}{GDP_{e-1}}\right)$$

Equation 2: We construct the one-term GDP growth rate, GDP_{1t}, where GDP_e is the GDP (in USD) during the election year *e*, *n* is the number of years since the last election, and GDP_{e-n} is the GDP during the election year prior to the election year *e*:

$$GDP_{1t} = ln\left(\frac{GDP_e}{GDP_{e-n}}\right)$$

Equation 3: We construct the one-year GNI growth rate, GNI_{1y}, where GNI_e is the GNI (in USD) during the election year *e* and GNI_{e-1} is the GNI during the year prior to the election year *e*:

$$GNI_{1y} = ln\left(\frac{GNI_e}{GNI_{e-1}}\right)$$

Equation 4: We construct the one-term GNI growth rate, GNI_{1t}, where GNI_e is the GNI (in USD) during the election year *e*, *n* is the number of years since the last election, and GNI_{e-n} is the GNI during the election year prior to the election year *e*:

$$GNI_{1t} = ln\left(\frac{GNI_e}{GNI_{e-n}}\right)$$

Equation 5: We construct the one-year GDP per capita growth rate, $GDPC_{1y}$, where $GDPC_e$ is the GDP per capita (in USD) during the election year *e* and $GDPC_{e-1}$ is the GDP per capita during the year prior to the election year *e*:

$$GDPC_{1y} = ln\left(\frac{GDPC_e}{GDPC_{e-1}}\right)$$

Equation 6: We construct the one-term GDP per capita growth rate, $GDPC_{1t}$, where $GDPC_e$ is the GDP per capita (in USD) during the election year *e*, *n* is the number of years since the last election, and $GDPC_{e-n}$ is the GDP per capita during the election year prior to the election year *e*:

$$GDPC_{1t} = ln\left(\frac{GDPC_e}{GDPC_{e-n}}\right)$$

Equation 7: We construct the one-year GNI per capita growth rate, $GNIC_{1y}$, where $GNIC_e$ is the GNI per capita (in USD) during the election year *e* and $GNIC_{e-1}$ is the GNI per capita during the year prior to the election year *e*:

$$GNIC_{1y} = ln\left(\frac{GNIC_e}{GNIC_{e-1}}\right)$$

Equation 8: We construct the one-term GNI per capita growth rate, $GNIC_{1t}$, where $GNIC_e$ is the GNI per capita (in USD) during the election year *e*, *n* is the number of years since the last election, and $GNIC_{e-n}$ is the GNI per capita during the election year prior to the election year *e*:

$$GNIC_{1t} = ln\left(\frac{GNIC_e}{GNIC_{e-n}}\right)$$

Equation 9: We construct the one-year unemployment growth rate, UE_{1y}, where UE_e is the unemployment rate during the election year *e* and UE_{e-1} is the unemployment rate during the year prior to the election year *e*:

$$UE_{1y} = ln\left(\frac{UE_e}{UE_{e-1}}\right)$$

Equation 10: We construct the one-term unemployment growth rate, UE_{1t}, where UE_e is the unemployment rate during the election year *e*, *n* is the number of years since the last election, and UE_{e-n} is the unemployment rate during the election year prior to the election year *e*:

$$UE_{1t} = ln\left(\frac{UE_e}{UE_{e-n}}\right)$$

Equation 11: We construct the one-year inflation growth rate, I_{1y}, where I_e is the inflation rate during the election year *e* and I_{e-1} is the inflation rate during the year prior to the election year *e*:

$$I_{1y} = \frac{I_e - I_{e-1}}{I_{e-1}}$$

Equation 12: We construct the one-term inflation growth rate, I_{1t}, where I_e is the inflation rate during the election year *e*, n is the number of years since the last election, and I_{e-n} is the inflation rate during the election year prior to the election year *e*:

$$I_{1t} = \frac{I_e - I_{e-n}}{I_{e-n}}$$

Equation 13: We construct the one-term voter turnout growth rate, V_{1t}, where V_e is the voter turnout rate during the election year *e*, n is the number of years since the last election, and V_{e-n} is the voter turnout rate during the election year prior to the election year *e*:

$$V_{1t} = ln\left(\frac{V_e}{V_{e-n}}\right)$$

Cleaning Data

The independent variables vary in historical availability. While all variables are available on an annual basis, data such as GDP, GDP per capita, inflation, and voting turnout have been collected since 1960 for all countries studied. However, the availability of GNI and unemployment rate are country-dependent. The Netherlands, Sweden, Australia, and Austria report their GNI starting in 1962, while the United Kingdom, New Zealand, and Denmark begin reporting GNI in the late 1960s and 1970s. Canada's GNI is only available starting in 1999. When GNI data is unavailable, the corresponding GDP growth rates is used in place of GNI growth rates. To be more specific, for the calculation of one-year GNI growth rate, if GNI data is unavailable for either the election year or the year prior, the one-year GDP growth rate is used instead. The same applies to one-term GNI growth rate, and one-year and one-term GNI per capita growth rates. Similarly, when the unemployment rate is unavailable, the nearest year's unemployment value is used.

A time sample from 1960 to 2021 is collected. However, since the construction of one-year and one-term growth rates require data from prior years to be available, data for the first year is unable to be calculated. Thus, all data from 1960 is eliminated. The final dataset consists of annual data from 1961-2021.

Variable Correlation

The inflation rate and one-year inflation growth rate consistently maintain strong correlations towards election outcomes, reinforcing the conclusion by Chappell and Veiga (2000) that voters react to increased inflation. In contrast, the one-term inflation growth rate was less significant. Of all the economic variables with one-year and one-term values, all one-year rates have a stronger, albeit negative, correlation with election outcomes than their one-term counterparts, suggesting that voters are more sensitive to recent economic changes.

Apart from economic factors, the one-term unemployment growth rate and one-term voter turnout growth rate show the most significant correlation with election outcomes. Factors collected from a single year, namely the unemployment rate and voter turnout rate in the year of the election, were insignificant.

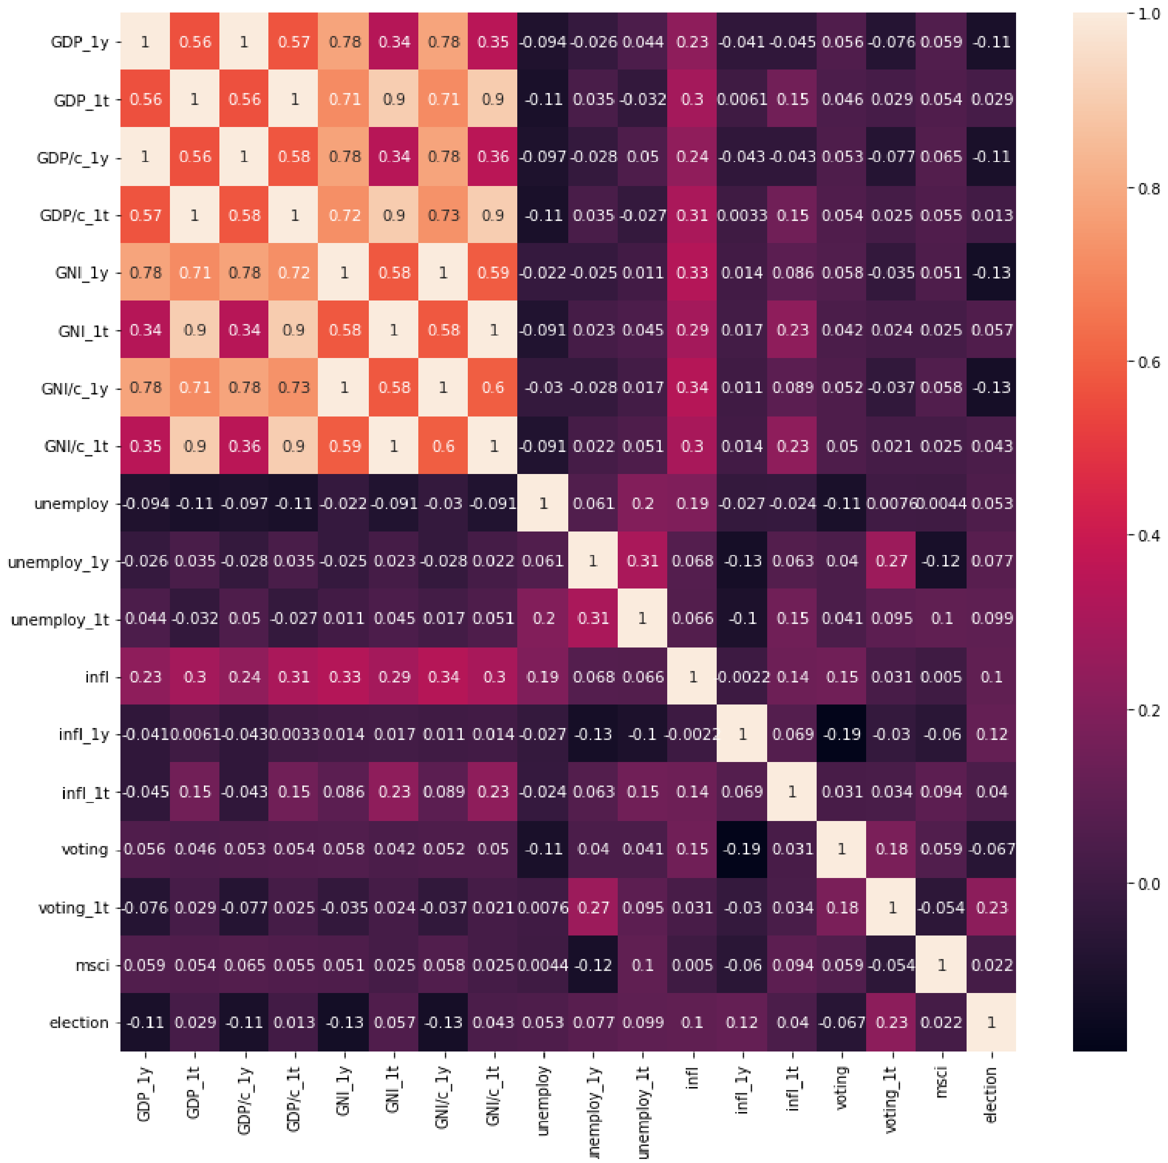

Figure 1. Heat map of all feature variable correlations.

Methodology

Software

Python 3.7 is used for all data processing and analysis. The dataset is cleaned and processed using Numpy and Pandas. Scikit-learn is utilized to implement the logistic regression, support vector machines, K-nearest neighbors, Gaussian process, Gaussian naïve Bayes, decision tree, random forest, AdaBoost, XGBoost, and multilayer perceptron classifiers.

Preprocessing

Data for the machine learning models is prepared by using 80% training data and 20% testing data. Feature variables are scaled to a standard distribution with a mean of 0 and variance of 1 using scikit-learn's Standard-Scaler.

Models

A myriad of machine learning classification models are utilized to predict party switches in federal elections of democratic countries. The models include logistic regression, support vector machines, K-nearest neighbors, Gaussian process classifier, Gaussian naïve Bayes, decision tree, random forest, AdaBoost, XGBoost, and multilayer perceptron. The listed models are applied to all eleven of our compiled national election datasets: the United States, Canada, the United Kingdom, the Netherlands, Austria, Norway, Sweden, Denmark, Australia, India, and New Zealand.

Results

Models

The multilayer perceptron with two hidden layers gives the most accurate classifications when forecasting party switches during democratic federal elections. On the contrary, Gaussian methods like Gaussian process classifier and Gaussian naive Bayes give the least accurate classifications. Table 2 lists a comprehensive overview of the performance of various models.

Table 2. Classifiers and their accuracy and variance for the prediction of party switches in democratic federal elections.

Model	Accuracy ± 95% Confidence Interval
Logistic regression	0.6644 ± 0.0308
Support vector machines	0.6343 ± 0.0278
K-nearest neighbors	0.5829 ± 0.0244
Gaussian process classifier	0.5886 ± 0.0265
Gaussian naive Bayes	0.5771 ± 0.0327
Decision tree	0.6343 ± 0.0353
Random forest	0.5886 ± 0.0282
AdaBoost	0.5657 ± 0.0381
XGBoost	0.6057 ± 0.0293
Multilayer perceptron	0.6968 ± 0.0201

Discussion

This research substantiates previous findings that economic and social forces play a significant role at the ballot box. Moreover, feature variable correlations show that one-year inflation growth rate, one-term unemployment growth rate, and one-term voter turnout growth rate influence election outcomes the most. Correlations across economic indicators consistently display that voters are more sensitive to short-term economic changes, specifically changes within one year of the election, compared to long-term changes over the term of the incumbent. This can be attributed to voters remembering recent changes to standards of living more vividly.

A major benefit of utilizing machine learning models is that prediction accuracy improves as more data is added. When the models are trained with data from only full democracies (countries that consistently score 8 or higher on the 10-point Democracy Index), which excludes the United States and India, all of the models predict party switches with less than 50% accuracy. However, with the addition of data from the United States and India, the multilayer perceptron and logistic regression classifiers in particular improved in overall accuracy as well as their precision and recall, resulting in more reliable predictions for all election outcomes. As more data becomes available over time, the models can be further optimized.

Conclusion

A wide array of machine learning models was used to predict the federal election outcomes of eleven democratic countries across the globe: the United States, Canada, the United Kingdom, the Netherlands, Austria, Norway, Sweden, Denmark, Australia, India, and New Zealand. First, the countries were selected based on the Democracy Index and population. Second, the most significant factors correlated with election outcomes were assessed. Third, a comparative analysis of machine learning models for the prediction of election outcomes was executed.

The results reinforce the conclusion made by Chappell and Veiga (2000) that voters react to increased inflation. Of all variables analyzed in the research, the one-year inflation growth rate, one-term unemployment growth rate, and one-term voter turnout growth rate show the strongest correlation toward election outcomes. Multilayer perceptron with two hidden layers provided the highest classification accuracy when forecasting whether the incumbent party stays in power. This is in contrast to Gaussian methods like Gaussian process classifier and Gaussian naive Bayes, which resulted in the worst performance out of the examined models.

Limitations

One major limitation was the data availability of variables, especially the unemployment rate. All eleven countries studied began reporting their unemployment rates in the 1990s, resulting in the training and test data showing a flat unemployment rate for the years from 1961 to 1990. Given the relatively higher correlation of one-term unemployment growth rate and election outcome, more unemployment data could yield even stronger results.

Additional features could also improve the predictions. Past literature showed promise in social sentiments collected from social media platforms, which were not included in this research. However, since federal elections generally occur once every three to five years, the number of features had to be limited to maintain a reasonable ratio between data points and features. Over time, as more elections occur, more data will become available, fostering the subsequent growth and advancement of literature. The models could also be applied to local elections.

Acknowledgments

The author especially thank Dr. Ganesh Mani from Carnegie Mellon University for his guidance, advice, and mentorship throughout the entirety of the research process. The author also thank Scholar Launch and Alfred Renaud for facilitating the relationship with Dr. Ganesh Mani and for their general support in the formulation of the research paper.

References

Budiharto, W., & Meiliana, M. (2018). Prediction and analysis of Indonesia Presidential election from Twitter using sentiment analysis. *J Big Data* 5, 51. https://doi.org/10.1186/s40537-018-0164-1

Chappell, H., & Veiga LG. (2000). Economics and elections in Western Europe: 1960–1997. *Electoral Stud.* 19. https://doi.org/10.1016/s0261-3794(99)00047-5

Kennedy, R., Wojcikand, S., & Lazer, D. (2017). Improving election prediction internationally. *Science,* Vol 355, Issue 6324, 515-520. http://doi.org/10.1126/science.aal2887

Lewis-Beck, M., & Stegmaier, M. (2000). Economic determinants of electoral outcomes. *Annual review of political science,* 3.1, 183-219. https://doi.org/10.1146/annurev.polisci.3.1.183

Nausheen, F., & Begum, S. (2018). Sentiment analysis to predict election results using Python. *2nd International Conference on Inventive Systems and Control (ICISC).* https://doi.org/10.1109/ICISC.2018.8399007

Paldam, M. (1991). How Robust is the Vote Function? A Study of Seventeen Countries over Four Decades. In H. Norpoth, M. S. Lewis-Beck , & J-D. Lafay (Eds.), Economics and Politics: The Calculus of Support (pp. 9-31). University of Michigan Press.

Tsai, M., Wang, Y., Kwak, M., & Rigole, N. (2019). A Machine Learning Based Strategy for Election Result Prediction. *International Conference on Computational Science and Computational Intelligence (CSCI)*, 1408-1410. https://doi.org/10.1109/CSCI49370.2019.00263