Mapping the Media in the Americas Project

by

Brett Wiens

AN EXEMPLAR ANALYSIS USING THE MAPPING THE MEDIA IN THE AMERICAS DATABASES
PREPARED FOR THE MEMBERS OF PARTICIPATING COUNTRIES
IN PARTIAL FULFILMENT OF THE HANDOVER OF THE WEBSITE DESIGN FILES, ARCIMS AND ARCGIS FILES

LATIN AMERICAN RESEARCH CENTRE
&
DEPARTMENT OF GEOGRAPHY
CALGARY, ALBERTA
APRIL 2008

Table of Contents

Introduction

In the most fundamental sense, democracy is a system of government by the people. Theoretically, every citizen of a democracy has an equal say in who leads the country, and what kind of policies govern their system. In practice, it is very difficult to establish perfect equality. In a representative democracy, such as found in Canada, inequalities of electoral divisions based on population distribution and the democratic system itself create electoral bias. Geographic information systems (GIS) can be used by democratic systems as a tool to determine what kind of people are voting for which policies, either as a tool by political parties as an effort to define party strategy to better represent their constituents or by government to try to ensure a fair allocation of seats.

There are a plethora of democratic systems in place in the world, and even more possibilities. Proportional representation guarantees that everyone's vote counts towards who gets elected, but proportionally elected representatives are not elected based on spatial boundaries, and the people are not represented by one person, and the representatives are not responsible to any set of constituents. Another form of government, used in Canada is the plurality-majority system. In this system, representatives are sent to parliament by the election of citizens from their particular spatial area, or riding. In a plurality majority system, a relatively small proportion of the popular vote is able to put a majority government in power. This is because smaller parties split the vote between them, and even with a minority of the vote, the winner is able to win the first-past-the-post system.

Sociodemographics are able to assist in the equitable distribution of political power throughout a country in any system. Through improved understanding of the electorate, the government is better able to respond to their needs. Furthermore, this kind of study can make finding the root of the issues easier. It is often more important to ask the question, why, as opposed to just what, when, where or whom. Policy is an effort to address how. It is difficult to prepare effective policy when the true problem or issue remains a mystery.

In an effort to analyze the Canadian electoral system, several studies were conducted of the 2000, 2004 and 2006 elections. This manual will explain the process used to analyze those elections and a brief summary of the results. Full results can be accessed from the University of Calgary. The objective of these studies was to determine what kind of people are more likely to vote for which party, based on their socioeconomic status.

CHAPTER ONE: Data

In many countries, it is possible to obtain specific sociodemographic data from several sources; these may be public corporations, such as Statistics Canada, or private corporations. Information may be obtained from the Mapping the Media in the Americas databases for high-level analyses. Analysis for Canadian elections has been completed using specialized variables believed to be important indicators of voting behaviour. A list of data available from the Mapping the Media in the Americas project is included in the MMA Training Manual on GIS. Data includes spatial data, electoral information, sociodemographic data and media data.

1.1 Acquire the data for your research

It should be noted that data varies from country to country, based on questions asked by the organization responsible for gathering census data. Some information that is very important in one country may prove to be unimportant in other regions. It is through an effort of trial and error, based upon literature and/or experience, that the best models are generated.

The Canadian Government gathers a very detailed, large quantity of data from the census. In total, there were 1,716 variables available for analysis in the 2001 census, the closest census to the 2000 election. It is best to correlate information from proximal time periods (2000 and 2001) than to try to estimate variability based on antiquated information. The 2000 election matched with information from the 1931 census would perhaps find interesting results for analysis, but would unlikely be useful for any policy or analysis. Many of the variables available are not particularly well suited to analysis in terms of political study. Some data is very specific, on country of origin or language, for example, and does not well represent the general population or their voting behaviour.

Fifty-three variables were selected based on research into similar analyses conducted by other academics. Some variables selected included:

Average income
Population with income > $60,000
Percent of income obtained from transfer payments
Population with less than a grade 9 education
Full-time income
Part-time income
Population with income < $30,000
Average family income
Population working in management fields
Population with university education
Population with less than a grade 13 education
Population educated in natural sciences
Population working in trades
Participation rate
Female participation rate
Male participation rate
Employment rate
Unemployment rate
Etc.

1.2 Select which of the available variables may be appropriate for your region

Many of these variables demonstrate similar information, or are likely to be highly correlated. Employment rate, participation rate and unemployment rate are strongly correlated, knowing any allows the calculation of the third easily. Similarly, areas with high average family incomes are also likely to have high personal average incomes. It is important that models do not use too many variables, especially those that are highly correlated, lest they over predict actual phenomena. One rule of thumb is to maintain no more than fifty more than eight times the number of cases. If there are 100 census divisions, the model may only make use of 6 independent variables. There will always be some level of correlation between variables, and it is critical to minimize intercorrelation of predictor variables in models.

The goal of any model is to explain the effect that independent variables have upon the dependent variable. The model must endeavour to maximize the correlation between the independent variables and the dependent variable, while minimizing the correlation between the independent variables. In order to achieve this, the fifty-three variables were fed into a statistics package called SPSS along with the dependant variables (the proportion of votes in each riding for each political party.) The resulting correlations were noted and were visualized in a correlation web. A correlation web shows the strength of correlations between variables. Variables with high correlation (90-100%) are linked by thick red lines, medium correlation (80-90%) by medium width, black lines, and low correlation (70-80%) by thin light blue lines. Some variables showed minimal correlation and were left unattached in the correlation web.

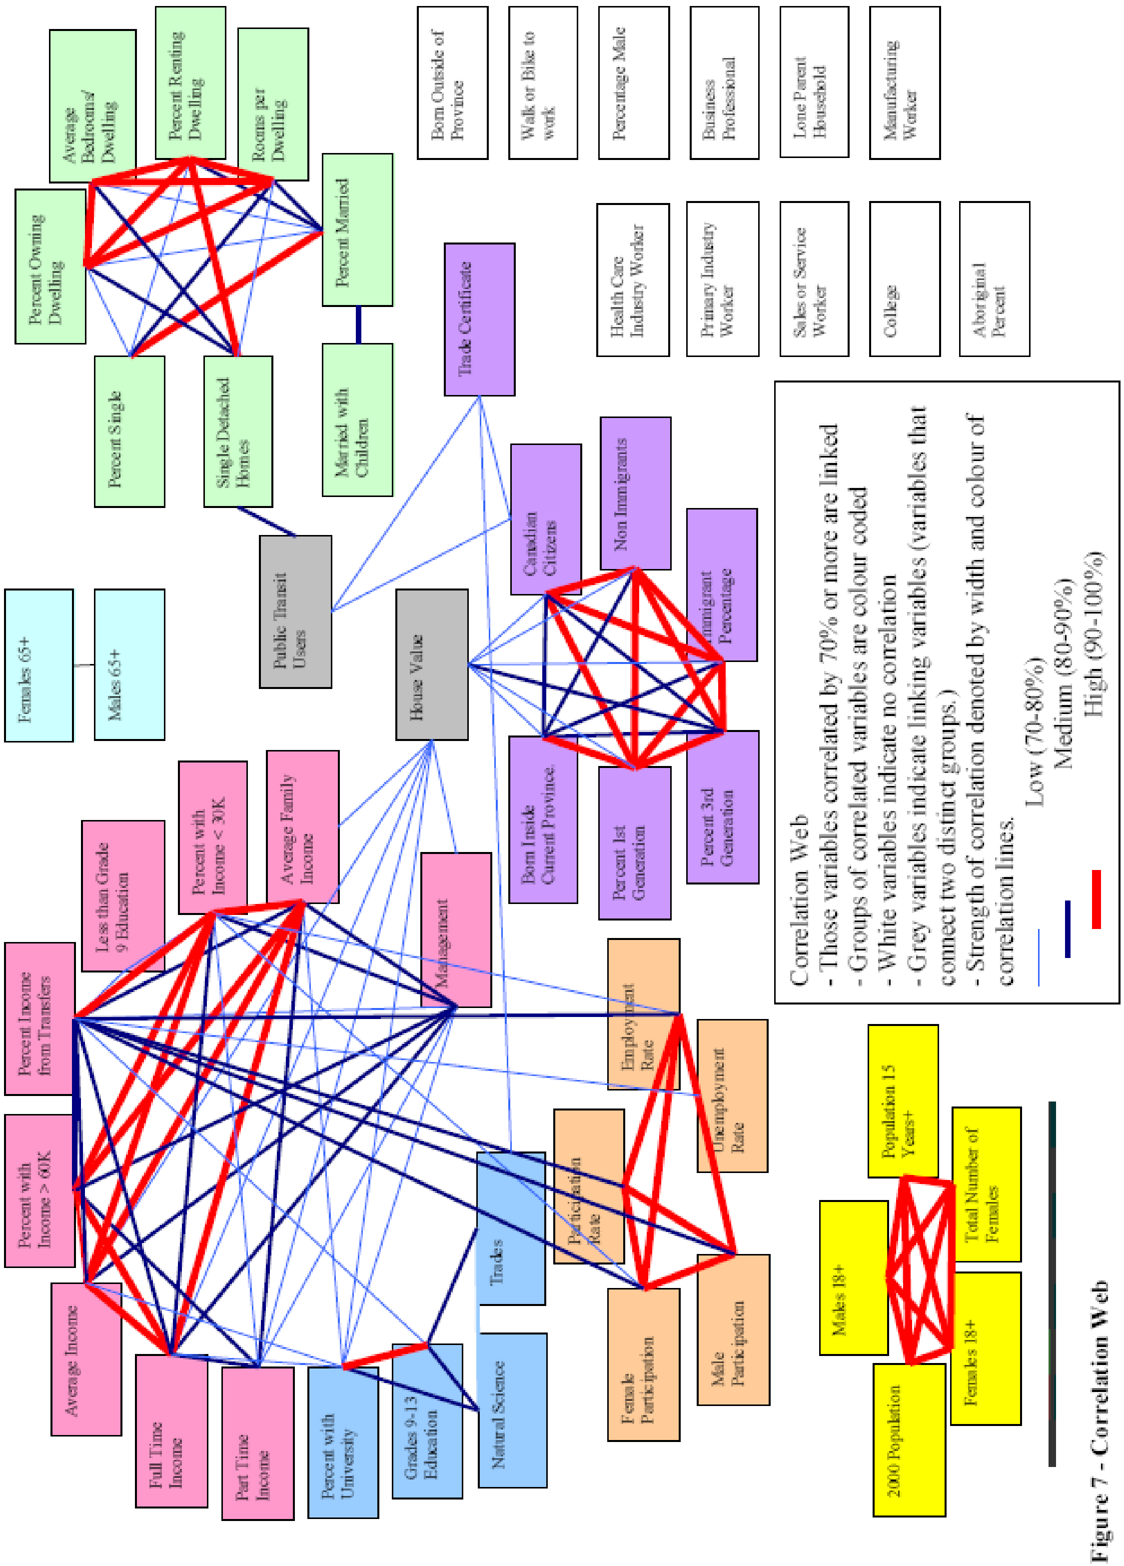

Figure 1: Correlation Web

1.3 Identify correlation between selected variables

By moving around the variables, groups of correlated variables can be created. In the course of this study, it was found that most groups are fairly easily identifiable based on the strongest correlations first. The groups that were established in this analysis were income, education, employment/participation, immigration, living situation and population. There are practically endless possibilities for variable selection, and limited time and money to do such analyses, so four case rules were created eliminate multicollinearity.

- The first case uses two rules:
  1. Maximize correlation between the independent and dependent variables.
  2. Maximize the number of variables.
- The second case inversed the rules:
  1. Maximize the number of variables
  2. Maximize the correlation between the independent and dependent variables
- The third case selected the ten independent variables most correlated to the dependent variables, then removed those that were correlated to one-another.
- The final case strictly maintained maximum independent-dependent correlation. The independent variable most highly correlated to the dependent variables was protected and any variable correlated with it was removed. The highest remaining independent variable was then protected and those correlated to it were removed, until no further multicollinearity existed.

Table 1: Correlation Cases

Case 1		Case 2	
Born Outside of Province	Unemployment Rate	Born Outside of Province	Part Time Income
Health Care Industry Worker	Percent with Income >60K	Health Care Industry Worker	Less than Grade 9 Education
Primary Industry Worker	Less than Grade 9 Education	Primary Industry Worker	Trade Certificate
Sales or Service Worker	Percent Married	Sales or Service Worker	Percent 1st Generation
College Degree	Percent in Single Detached Homes	College Degree	Percent 3rd Generation
Walk or Bike to Work	2000 Population	Walk or Bike to Work	Percent with Income >60K
Percentage Male	Female Participation Rate	Percentage Male	Natural Science
Business Professional	Percent 3rd Generation +	Business Professional	Participation Rate
Lone Parent Household	Females 65+	Lone Parent Household	Unemployment Rate
Manufacturing Worker	Natural Science	Manufacturing Worker	Males 65+
Aboriginal Percent		Aboriginal Percent	2000 Population
		Percent Single	Public Transit Users
		Percent Married with Children	
Case 3		**Case 4**	
Percent 3rd Generation	Percent in Single Detached Homes	Born Outside of Province	2000 Population
Female Participation	Manufacturing Workers	Health Care Industry Worker	Females 65+
Less than Grade 9 Education	Primary Industry Workers	Primary Industry Worker	Grades 9-13 Education
Percent Married	Born Outside of Province	Sales or Service Worker	Percent 3rd Generation
		College Degree	Percent with Income >60K
		Walk or Bike to Work	Female Participation
		Percentage Male	Unemployment Rate
		Business Professional	Trade Certificate
		Lone Parent Household	Less than Grade 9 Education
		Manufacturing Worker	Percent Married
		Aboriginal Percent	Single Detached Homes

These four cases were used as the basis for the first part of the regression process.

CHAPTER TWO: Proportionalising Variables

Many variables are presented from the census as raw numbers. There may be 51 people in an electoral district with university educations and 102 in another, based on this information, the model may assume that the second area has twice the concentration of university educated people. This is not the case however, if the first electoral district contains only 100 people, while the second represents 400, in reality, the second electoral district would have half the concentration of university educated persons when compared to the first. It is therefore crucial to convert the data from the census from raw data to proportions, in order to fairly compare one electoral district to another. In reality, this degree of malapportionment is unlikely to occur, but to some extent, unless each electoral district has exactly the same number of people, it will still happen.

To convert data into proportions, one extra piece of information is required, the total population in each electoral district. Fortunately, the original purpose of the census was to count the number of people for tax reasons. This data is often the most easily accessible, and is available from the Mapping the Media in the Americas Project.

To make the raw data into proportions, simply divide the raw data by the total population available to a certain piece of data. To calculate the population with less than a grade nine education, one would not include children, so divide the raw population with less than a grade nine education by the adult population. To calculate the proportion of homes with a certain floor material, divide the number of homes with that flooring material by the total number of homes.

2.1 Convert raw data into proportional data as required/possible

Table 2: Proportion conversions

Dissemination Area	Population	Families	Lone Parent	Less Than Grade 9	Proportion Lone Parent	Proportion Less Than Grade 9
126564	1890	700	140	268	0.2	0.141798942
351654	2678	956	236	359	0.246861925	0.134055265
654322	2415	1012	25	372	0.024703557	0.154037267
675649	1653	786	30	202	0.038167939	0.122202057
126315	1268	862	35	242	0.040603248	0.190851735
645844	2153	432	49	423	0.113425926	0.196470042
214589	2210	769	68	401	0.088426528	0.181447964
136774	1568	758	70	203	0.092348285	0.129464286
654353	1697	821	100	235	0.12180268	0.13847967
897623	1921	842	68	210	0.080760095	0.109318064
684684	1903	563	89	151	0.158081705	0.079348397

CHAPTER THREE: Regression Algorithms

One part of this analysis was to compare simple linear regression techniques. Linear regression includes a number of similar, but subtly different methods of model building and regression. The best algorithm is the one that explains the greatest amount of variance in the dependent variables, measured by the R^2 statistic.

There are three main types of linear regression, named based on the way the variables are added/removed from the model. The first type, enter, effectively throws every variable into the model. The second, forward stepwise, seeks to add variables one-by-one, starting with the independent variable most highly correlated to the dependent variable, based upon whether they improve the model significantly. (Significance is established as having a default entry F-statistic of 0.05). If a variable does not significantly improve the model, it is not added. The third type is backwards stepwise regression. In backwards stepwise regression, all variables start in the model and the variables that do not benefit the model, are removed, based on a removal F-statistic of 0.10. The smallest number of variables that predicts the greatest amount of variability is used in the final models.

For this research, each of the three regression techniques were run for each of the five major Canadian parties that participated in the federal election in 2000 and for each of the variable cases. In total, 60 regressions were run, simultaneously comparing the four case rules of variables and the three regression techniques. It was found that the four cases and the three regression techniques were able to predict almost 50% (an R^2 value of 0.48). The backwards stepwise regression explained the greatest proportion of the variability when combined with variable selection case 2. Therefore, further analysis, and the geographically weighted regression would use the variables in case 2.

3.1 Select which linear regression technique will be used

(SPSS – Statistics Package for the Social Sciences - Used in Sample)

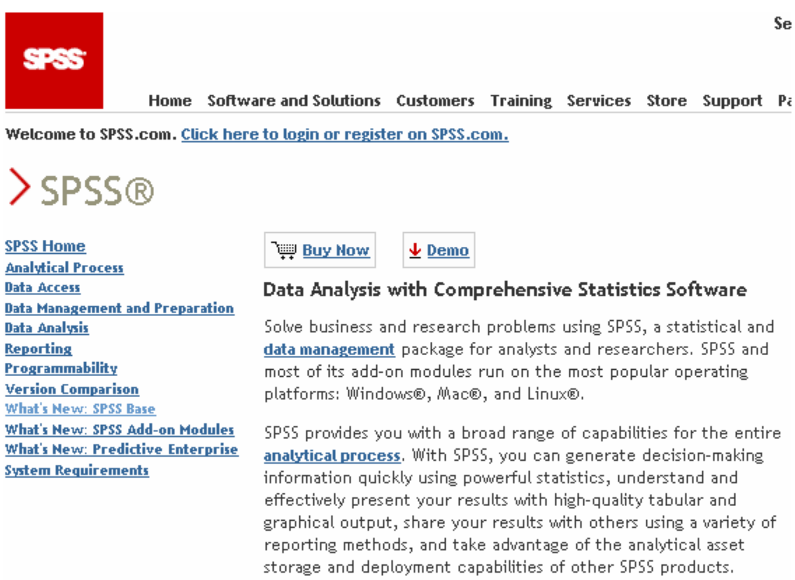

Figure 2: http://www.spss.com/SPSS/

3.2 Run a Linear Regression

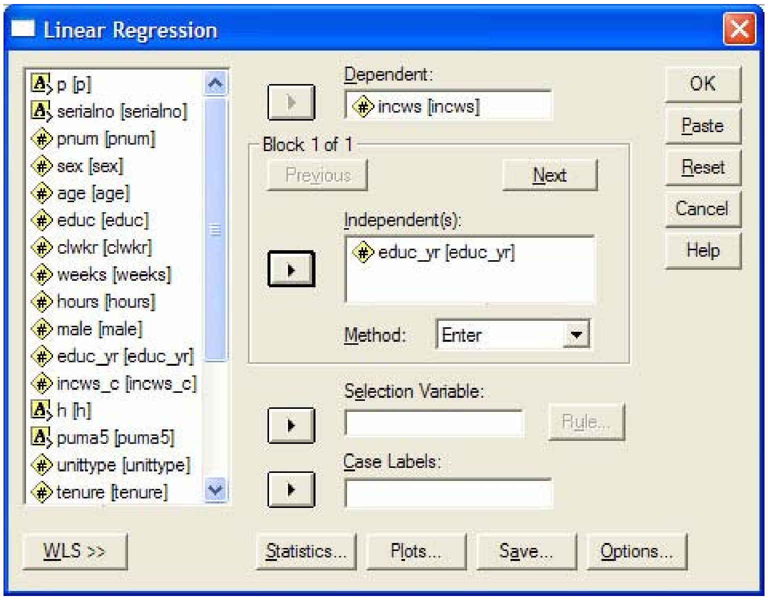

Figure 3: Regression variable selection

CHAPTER FOUR: Geographically Weighted Regression

Geographically weighted regression is a more recent statistical technique. While linear regression tends to assume that trends are effectively independent of location, GIS, and the inherently spatial distribution of voting patterns declares that this is an overly simplistic way to look at the world. Geographically Weighted Regression (GWR) takes into account spatial non-stationarity by calculating relationships between independent and dependent variables on a local, rather than purely global scale. More information on Geographically Weighted Regression, including the formulae and software can be found on Dr. Fotheringham's website: http://ncg.nuim.ie/ncg/GWR/index.htm.

GWR is useful in accounting for Tobler's first law of geography; those areas closer to an unknown location will have a greater influence than those areas further away. As a result, a bandwidth or kernel must be selected to define the weight of proximal information. There are two types of kernel that may be selected, a fixed kernel, which maintains a distance measure, no matter how many neighbours fall within a distance of a point, they all exert an influence, and the second is an adaptive kernel. An adaptive kernel chooses the number of neighbours nearest, unto the kernel number, to exert the influence. A fixed kernel is best for regularly spaced phenomena, if a bandwidth of 100m is selected and no neighbours occur within that area, then it is predicted only by its own data. Where distribution is irregular, as was the case in Canada, where the overwhelming majority live in the South, an adaptive kernel is best, though neighbours further away will still exert minimal affects on the results.

A kernel may be specified by the user, or it may be calculated using Akaike's Information Criterion. AIC = -2loglike+2K (K = number of parameters to be estimated). The best model is the one with the smallest sample-size-adjusted AIC value. The program creates models comparing the AICc values until the lowest AICc possible is found.

4.1 Select your kernel, or allow the program to automatically calculate an appropriate kernel. Record the kernel for analysis purposes.

Unlike linear regression techniques, geographically weighted regression does not produce a single formula which may explain any point at any location; however, it can calculate a specific formula for any unknown point. Each formula is unique at any location in the world.

4.2 Run a Geographically Weighted Regression

So once the data has been reformatted into an excel file with rows and columns, a number of steps need to be carried out before it can be mapped. Each row is a record of data for a certain geographical unit, and each column is a certain characteristic or attribute inherent to that particular geographical unit. Since each geographical unit is unique within that country, it needs to be assigned a unique identification number. This is important for later steps as these databases need to be spatially joined to the geographical boundary data in order for a map to be created.

CHAPTER FIVE: Linear Regression Results

Linear regression produces a number of statistics; a few of the most useful statistics are the standardized linear regression coefficients. These numbers, for which there is one for every independent variable, are actually the blueprint for calculating the dependent variable. The dependent variable can be calculated for any census area using the following formula.

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 X_{i\ldots}$$

Y = Dependent Variable
β= Standardized Regression Coefficient
X = Independent Variable
i = Current Location

By calculating the dependent variable using this formula, the model is able to predict as much of the vote distribution as possible from the independent variables used. The model can be checked for accuracy by comparing the actual vote proportion to those predicted by the model. The difference between the predicted and the actual vote totals is the residual. If you sum the squared values of the residuals, and add the summed squared regression values and the total by the summed squared regression values, you get an R^2 statistic, which is the predictive value of the model. An R^2 of 0.8 would indicate that the model predicts 80% of the variance in the dependent variable. This statistic is generated automatically in most software packages, including SPSS.

The sample from this study calculated linear regression for all five parties, for simplicity sake, only two, the two largest parties, the Liberal and Canadian Alliance, will be discussed in detail.

5.1 Determine the predictive formula for each party using the standardized regression values. Beta values from the statistical program.

Coefficients(a)						
Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
8	(Constant)	33.800	14.659		2.306	.022
	tradecert	-.620	.297	-.196	-2.089	.038
	ltgrade9	-.665	.174	-.372	-3.830	.000
	perc3rdgen	.437	.061	.824	7.137	.000
	over60k	-.353	.186	-.171	-1.903	.058
	pop2000div	-.263	.057	-.456	-4.582	.000
	publictran	.181	.096	.197	1.885	.060
	bornoutpro	.185	.072	.178	2.560	.011
	singleperc	-.594	.144	-.406	-4.125	.000
	marrwchild	.817	.125	.671	6.558	.000
	aboriginal	-.376	.093	-.331	-4.037	.000
	health	-2.020	.521	-.224	-3.877	.000
	walkorbike	.715	.168	.409	4.266	.000
	business	-.599	.253	-.219	-2.367	.019
	primaryind	-.486	.152	-.264	-3.205	.002
	loneparent	.004	.001	.421	4.001	.000
	college	-.609	.218	-.203	-2.794	.006
a Dependent Variable: progressiv						

Figure 4: Coefficients

5.2 Observe and record the R-square value.

The regression coefficients for the Liberal Party showed that fifteen variables were important for predicting Liberal votes, either in a positive or negative way. The Liberal vote proportion is defined by the following formula:

> Liberal = 68.073 – 1.775 * (Percent with trade certificates) + 0.217 * (Percent 3rd generation or more) + 1.078 * (Unemployment rate) + 0.003 * (Percent males aged 65 or older) – 0.000347 * (Population) – 0.689 * (Percent born outside their province) – 0.41 * (Percent single) + 0.512 * (Percent married with children) – 3.614 * (Percent working in health care) + 0.655 * (Percent who commute by bike or on foot) – 0.478 * (Percent working in business fields) – 1.00 * (Percent working in primary industry) + 0.003 * (Percent lone parents) – 0.993 * (Percent working in sales or service industries) + 1.484 * (Percent with college degrees)

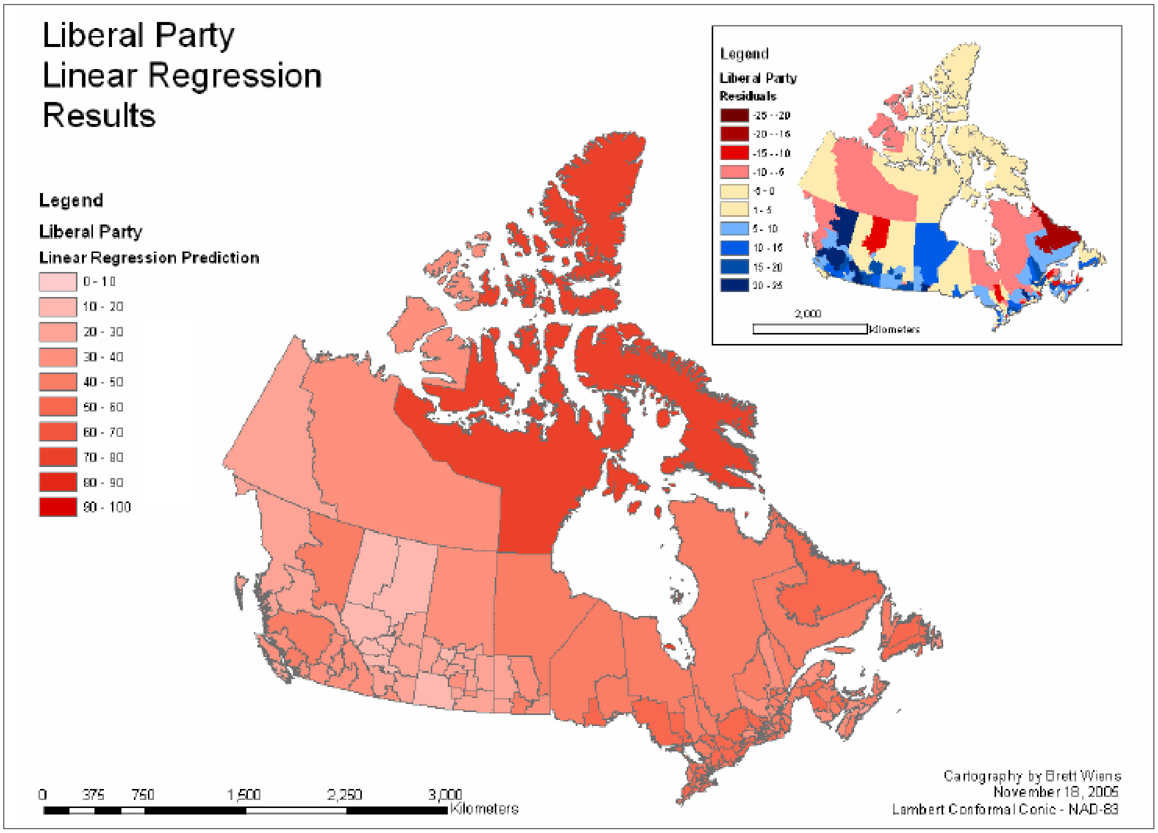

Figure 5: Liberal Party Linear Regression Results

The regression coefficients for the Canadian Alliance (CA) used fourteen variables to predict their vote proportion. The CA vote proportion is defined by the following formula:

> Canadian Alliance = 49.2 + 1.132 * (Percent with trade certificates) – 0.614 * (Percent 3rd generation or more) – 0.897 * (Percent earning more than $60,000) - 0.79 * (Unemployment rate) – 0.003 * (Percent males aged 65 or older) + 0.00043 * (Population) + 0.83 * (Percent born outside their province) – 0.258 * (Percent married with children) +1.923 * (Percent working in health care) – 0.304 * (Percent who commute by bike or on foot) – 0.992 * (Percent working in business fields) + 0.96 * (People working in primary industry) – 0.002 * (Percent lone parents) – 1.052 * (Percent working in manufacturing)

This shows us a few of the typical differences between people who vote for the Liberal party and those who vote for the Canadian Alliance. Generally, as the two largest parties, the type of jobs, families, education, etc. that predict a positive vote for one party, directly predict a negative correlation to the other party. The Liberals are popular areas with people whose families have lived in Canada for a long time, and do better in areas where there is higher unemployment. The Alliance did well in areas where people worked in trades and in electoral districts with higher populations (typically in Canada, cities are underrepresented proportionally).

When reading the formulae which result from this analysis, the most basic interpretation is that those variables with a negative beta value are those that have a negative impact on the dependent variable, or the party's vote proportion. Those which have a positive beta value are positive indicators of the party's success.

5.3 Determine which variables are positive and negative indicators.

There is a major caveat to this kind of analysis. It may appear that people with trade certificates are voting strongly for the Canadian Alliance, however, this is only one possibility. It is also possible that in areas where there are a lot of tradespeople the people working in trades are voting against the Canadian Alliance or splitting their vote among parties, while their neighbours without trade certificates are voting strongly for the Alliance. Short of asking every person in a country about their census information and their political leanings, it is not possible to infer more than this level of detail.

Linear regression was able to predict between 31.4% of the vote for the New Democratic Party, 38.5% for the Progressive Conservatives, 49.9% for the Liberals, 69.6% for the Canadian Alliance, and 80.4% for the Bloc Quebecois. Canadian politics are very regional in nature, with the country divisible among a few very distinct regions: Western Canada, Ontario, Quebec, the Maritime Provinces and the North. In the 2000 election, the West was won by the Canadian Alliance, Quebec was won by the Bloc Quebecois, and the remainder went to the Liberals. This pattern explains many of the failures of the linear regression in predictive value.

CHAPTER SIX: Geographically Weighted Regression Results

As discussed, GWR does not provide a simple formula output, but rather unique calculations for every point based on its neighbours. As a result, the typical output from the geographically weighted regression is a map. The software directly outputs an ESRI-compatible .e00 format for both the regression and residuals, as well as a list of statistics.

```
*************************************************
*       Geographically Weighted Regression      *
*               Release 3.0.1                   *
*            Dated: 06-vii-2003                 *
*       Martin Charlton, Chris Brunsdon         *
*            Stewart Fotheringham               *
*     (c) University of Newcastle upon Tyne     *
*************************************************
Program starts at: Wed Nov 16 14:37:49 2005

** Program limits:
** Maximum number of variables.....      52
** Maximum number of observations..  80000
** Maximum number of fit locations.  80000

** Variables in the data file...
 ID       Xcoord   Ycoord   POP2000  MALEPERC MALE65PL SINGLEPE MARRWCHI
 LONEPARE BORNOUTP PERC3RDG ABORIGIN PARTICIP UNEMPLOY BUSINESS NATURALS
 HEALTH   SALESSER PRIMARYI MANUFACT PUBLICTR WALKORBI LTGRADE9 TRADECER
 COLLEGE  PARTIMEI OVER60K  LIBERAL  PROGRESS N_D_P_   CANADIAN BLOC_QU╔

** Dependent (y) variable..........LIBERAL
** Easting  (x-coord) variable.....Xcoord
** Northing (y-coord) variable.....Ycoord
** No weight variable specified
** Independent variables in your model...
 MALEPERC MALE65PL SINGLEPE MARRWCHI LONEPARE BORNOUTP PERC3RDG ABORIGIN
 PARTICIP UNEMPLOY BUSINESS NATURALS HEALTH   SALESSER PRIMARYI MANUFACT
 PUBLICTR WALKORBI LTGRADE9 TRADECER COLLEGE  PARTIMEI OVER60K
** Kernel type:  Adaptive
** Kernel shape: Bi-Square
** Bandwidth selection by AICc minimisation
** Use all regression points
** Calibration history requested
** Prediction report requested
** Output estimates to be written to .e00 file
** Monte Carlo significance tests for spatial variation
** Casewise diagnostics to be printed

*** Analysis method ***
*** Geographically weighted multiple regression
** Cartesian coordinates: Euclidean Distance
**************************************************************
*        GEOGRAPHICALLY WEIGHTED GAUSSIAN REGRESSION         *
**************************************************************
 Number of data cases read: 301
 Observation points read...

 Dependent mean=  41.1328888
 Number of observations, nobs= 301
 Number of predictors,   nvar= 23
 Observation Easting extent:   4928551.5
 Observation Northing extent:  3121238.5
*Finding bandwidth...
  ... using all regression points
 This can take some time...
 *Calibration will be based on  301 cases
*Adaptive kernel sample size limits:      15    301
 *AICc minimisation begins...
           Bandwidth                     AICc
```

Figure 6: GWR Output

By incorporating the spatial component of voting behaviour in Canada, the GWR increased the prediction capability by as much as 20% for the Progressive Conservative Party. The model was able, for the Bloc Québécois, which only runs candidates in Quebec, was able to predict 85.4% of the vote using independent variables and spatial correlation. GWR was able to predict very accurately those regions where parties receive very few or no votes.. Its weakest areas are those where there are sharp borders between similar communities, such as the border between Québec and New Brunswick, in which GWR predicted some votes for a party that doesn't run any candidates in that province.

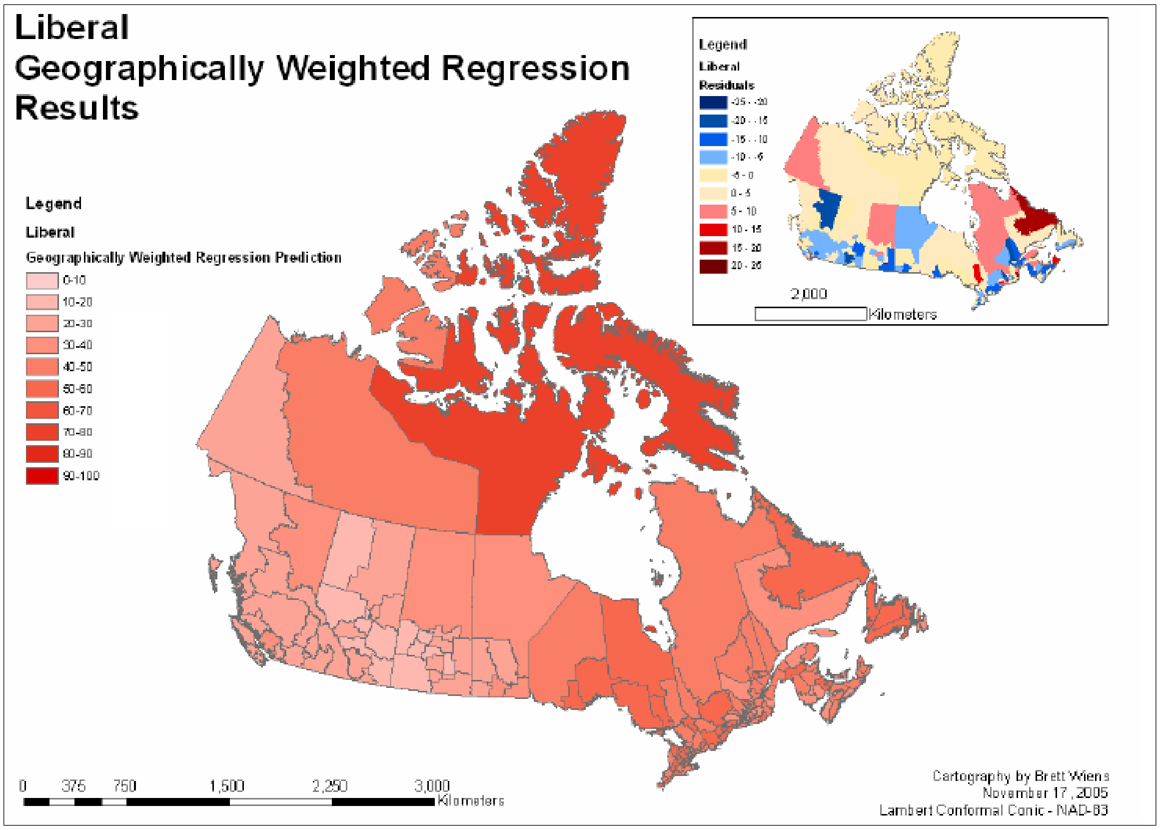

Figure 7: Liberal Party GWR Results

Another output of the geographically weighted regression is the kernel itself. Since the program calculated the appropriate kernel for each party, some information can be gleaned from that information alone. In the case of Canada, three parties had very large kernels, these were the parties that got relatively high or at least constant votes throughout the country, each of the Liberals, NDP and Canadian Alliance had bandwidths of over 200. The Bloc Québécois had a much smaller bandwidth, taking into account only 140 electoral districts. This shows that the Bloc was the most regional party, the chances of the party winning were only high near other areas where they won, parties like the Liberal and Canadian Alliance were less effected by geography, gaining somewhat consistent results throughout the country.

CHAPTER SEVEN: Discussion

Using sociodemographic analysis with respect to election results, it is possible to explain the distribution of voting behaviour. This information can be used for a variety of purposes as discussed earlier. Ideally for the betterment of the democratic process, to better represent the people and to better serve their needs. Using linear regression techniques it is possible to see what kinds of people are more likely to vote for which parties, and hopefully why, and adjust policy to better serve their needs. For purely predictive purposes, a geographically weighted regression is likely to produce higher quality results. The subject matter presented here by no means represents the limits to this kind of research. The Mapping the Media in the Americas project has provided more data than just census information and electoral results. Information about media outlets and media representation can be used to analyze the effects of media on elections or on the population they serve. The information need not be limited to electoral or political purposes as delineated here, the methodology has practical applications in a wide variety of fields.

References

Fotheringham, S.A., Brundson, C., and Charlton, M. (2002) *Geographically Weighted Regression*. Chichester, WS: John Wiley & Sons Ltd.

Johnston, R.J. (2002) The 2000 Annual Political Geography Lecture: Manipulating maps and winning elections: measuring the impact of malapportionment and gerrymandering. *Political Geography,* 21. 1-31.

McGrew, J.C., and Monroe, C.B. (2000) *An Introduction to Statistical Problem Solving in Geography, Second Edition.* Boston, MA: The McGraw-Hill Companies, Inc.

Rogerson, P.A. (2001) *Statistical Methods for Geography*. London: Sage Publications Ltd.

Wiens, Brett (2006) *Analyzing the 2000 Canadian Federal Election*. Calgary, AB.