Applying Locally Weighted Scatterplot Smoothing to Spatial Time-Series Analyses of Metropolitan Regions

Steven Pham

NCRP Research Paper 245

February 2021

Funded by the Social Sciences and Humanities Research Council of Canada
through the Neighbourhood Change Research Partnership

www.neighbourhoodchange.ca

ISSN 0316-0068; ISBN 978-0-7727-2503-5

Social Sciences and Humanities Research Council of Canada

Conseil de recherches en sciences humaines du Canada

Canada

Applying Locally Weighted Scatterplot Smoothing to Spatial Time-Series Analyses of Metropolitan Regions

Steven Pham
February 2021, vi, 29 pp.
ISSN 0316-0068
ISBN 978-0-7727-2503-5

Neighbourhood Change Research Partnership
Factor-Inwentash Faculty of Social Work
University of Toronto
248 Bloor Street West
Toronto, Ontario, Canada
M5S1V4
E-mail: neighbourhood.change@utoronto.ca
Website: http://neighbourhoodchange.ca

This study and its dissemination are supported by grants from the Social Sciences and Humanities Research Council of Canada, which has funded a Faculty Insight Grant held by Dr. Alan Walks, and the Neighbourhood Change Research Partnership based at the University of Toronto (J. David Hulchanski, Principal Investigator).

NEIGHBOURHOOD CHANGE | Research Partnership
Trends | Processes | Consequences | Policy Interventions
www.neighbourhoodchange.ca

Executive Summary

Geography and planning frequently employ thematic mapping to provide visual summaries of complex socio-spatial phenomena that manifest across urban centres over defined periods of time. However, thematic mapping imposes limitations in terms of how data is graphically summarized, leading to difficulty of interpretation and approximation of changes or dynamics, as well as author bias. This paper proposes non-parametric data smoothing techniques, in particular Locally-Weighted Scatterplot Smoothing (LOESS), as an alternative method to traditional mapping techniques for visualising spatio-temporal dynamics and trajectories across metropolitan regions. This paper focuses on the research applications and utility of LOESS in graphically summarizing relationships between social, economic, and political variables, and their changes over time in an urban context. LOESS is compared to another non-parametric smoother, Exponential Smoothing, to fully demonstrate its advantages and disadvantages, followed by a discussion regarding the accuracy of LOESS. Ultimately, it is demonstrated that LOESS more accurately and clearly delineates spatio-temporal trends in comparison to Exponential Smoothing.

Authors

Steven Pham is a Graduate from the Masters of Science in Planning (MScPl) program at the University of Toronto. Previously, Steven completed undergraduate studies at the University of Toronto in Mississauga, having majored in Human Geography and GIS and graduating with high distinction. During his time at UofT, Steven completed an undergraduate thesis which delineated and defined neighbourhood change in Toronto, and a master's thesis which assessed attitudes towards suburban retrofitting in Mississauga. His undergraduate thesis was the Award Winner for Best Thesis in Geography, Environment or GIS Programs in 2017, and is currently being finalized for publication. Steven has presented his work as a research assistant at many geographical and planning conferences, including the American Association of Geographers and the Ontario Professional Planners Institute. His research interests include gentrification, socio-spatial inequality, neighbourhood change, and the geography of debt and finance. Currently, Steven is working as a land use planner at a GTA-based consulting firm providing professional planning opinion to developers.

Acknowledgements

The author would like to thank Dr. Alan Walks of the University of Toronto Mississauga for supervising the research project and for providing feedback on this technical paper. This research was supported by the Social Sciences and Humanities Research Council of Canada, through a faculty Insight grant to Dr. Walks. The author has no conflicts of interest with respect to the methods compared herein.

Table of Contents

List of Figures

1. Introduction

Visually representing social, economic, and political phenomena across time and space has been a key tool for social scientists in analyzing and documenting change at various scales, with the goal of ultimately affecting policy changes for public and private actors. Social scientists traditionally employ thematic maps to portray patterns and trajectories across urban areas, particularly over metropolitan regions.

However, while thematic maps are powerful tools for visually representing social, economic, and political phenomena across vast distances, their abilities to convey changes over time are more limited. Thematic maps categorize and classify data to represent and emphasize key spatial patterns, yet the necessity of data classification also hinders accurate and simple time-series analyses.

An alternative to thematic mapping involves multi-dimensional time-series graphing.Two-dimensional (2-D) and three-dimensional (3-D) time-series graphs can be adapted to summarize spatial phenomena, including summarizing changes in their distribution over time, often with greater detail and allowing for easy calculation of summary metrics. However, as anyone who maps spatial phenomena is aware, the spatial distribution of most variables is highly uneven, particularly over time, making it necessary to employ data smoothing techniques before graphing spatial data, thereby making them legible and meaningful.

Using non-parametric data smoothing techniques, spatial phenomena, including those for metropolitan dynamics and time-series trajectories, can be accurately graphically summarized. Non-parametric smoothing methods empirically obtain and represent the structure of the underlying data without constraint, and the flexibility of non-parametric methods allows them to reveal more complex relationships or processes that may be obscured by the use of more traditional parametric statistical methods, or those for which no theoretical models exist.

However, there are alternative data smoothing methods, and to date insufficient research has been conducted to analyze which might be most useful for representing spatial data, particularly the sort that varies based on distance from the centre of other established nodes. This paper considers two common data-smoothing alternatives used for graphically visualizing data trajectories: exponential smoothing, and locally-weighted scatterplot smoothing (LOESS). The objective of the paper is to compare the capabilities of these two options for graphically summarizing spatial phenomenon, such as metropolitan spatial time-series dynamics and trajectories where key social variables change in irregular fashion with distance from certain locations, such as Central Business Districts (CBDs), over time. In doing so, this paper both proposes and justifies an alternative method to traditional thematic mapping in conducting spatial time-series analysis of metropolitan regions.

The paper begins by examining the problem of representing spatial patterns in time-series in metropolitan areas, and the use of scatterplot smoothing as a solution. It then compares two key methods that have been developed in the literature for scatterplot smoothing in general, but whose different strengths and weaknesses have yet to be analyzed in relation to their application in spatial time-series data. Based on my analysis, I make an argument for selecting LOESS for scatterplot smoothing. In doing so, I outline the steps for calculating and applying LOESS curves, as well as its limitations. Finally, this paper concludes with a discussion of the results and implications for further refinement of the use of LOESS for future spatial time-series analyses of metropolitan regions. For ease of calculation, the analysis here is here restricted to two-dimensional (2-D) time-series graphing, but the same principles would be applicable to three-dimensional data series.

2. Capturing Metropolitan Dynamics and Trajectories

Metropolitan regions across the globe have experienced dramatic transformations, not least among inner and suburban zones (the inner-city, inner-suburbs, outer-suburbs, exurbs, etc.), and these have only increased in the post–Second World War era. In transitioning from Fordist to Post-Fordist and neoliberal policy regimes, metropolitan regions have seen waves of prosperity and decline emerge within and sweep across zones. The geographical literature is filled with scholarly works detailing the spatially-uneven processes driving prosperity and decline, such as white flight, suburbanization, deindustrialization, professionalization/tertiarisation, gentrification, (reverse) filtering, etc.

At a finer scale, all these processes have tremendously impacted local neighbourhoods and have created and deepened intra-zonal inequalities across metropolitan regions (see, for example, Walks on Canadian cities 2001, 2011). Traditionally, visually summarizing processes of prosperity and decline (and spatial data in general) across space and time has been achieved through mapping the metropolitan region at varying scales. However, despite the clear benefit of maps for pinpointing spatial locations of various phenomena, the effectiveness of monitoring changes in time-series census data via mapping also has its drawbacks. First, as is customary of thematic mapping, maps of socio-economic variables generally classify them within a series of colour or size-coded ranges, resulting in potentially inaccurate and inadequate approximations of changes across time, particularly given that a single colour (or other symbol/size) typically corresponds to a large range of values.

As well, thematic maps must be scrutinized carefully, as the method of data classification may bias visual perceptions and analyses depending on the creator's intentions or prejudices, with the size of the spatial units being mapped potentially misleading readers to the true underlying distribution (Monmonier 1995, pp. 40–42).Second, time-series analysis of spatial census data is often constrained by changing boundaries of census units if the boundaries are not standardized. For most of the datasets used by geographers to analyze metropolitan change, the changing boundaries of the spatial units being mapped can prevent an accurate analysis of in-situ decline or increase. Third, at finer scales often the sheer number of census units, combined with the complex and messy reality of spatial distributions of socio-economic variables, makes visual interpretation of changes across time daunting and laborious, making it necessary to develop single metrics or more simple graphing procedures to truly understand the degree and location of underlying changes.

Two-dimensional (2-D) time-series graphs can facilitate the monitoring of changes across time. By plotting the same dependent and independent variables at different points, the user is able to distinctly note changes in trends (if any) and visually approximate the degree of change between each time period. As graphs do not require data classification, they have an advantage over thematic maps in facilitating highly accurate readings. As well, time-series graphs typically involve fitted curves which summarize the relationship between variables, eliminating potentially laborious visual interpretation of many census units.

Of course, for time-series graphs a key independent variable will be time. If one is primarily interested in analyzing spatial phenomena over time, combining geographic analysis with the clarity offered by time-series graphs can provide insightful information that maps cannot. Since time-series curves are not constrained by census unit boundaries, it becomes possible to analyze changes across spatial units over time despite using discrete unstandardized boundaries.

Figure 1: University Education by Census Tract across the Toronto CMA, 1981

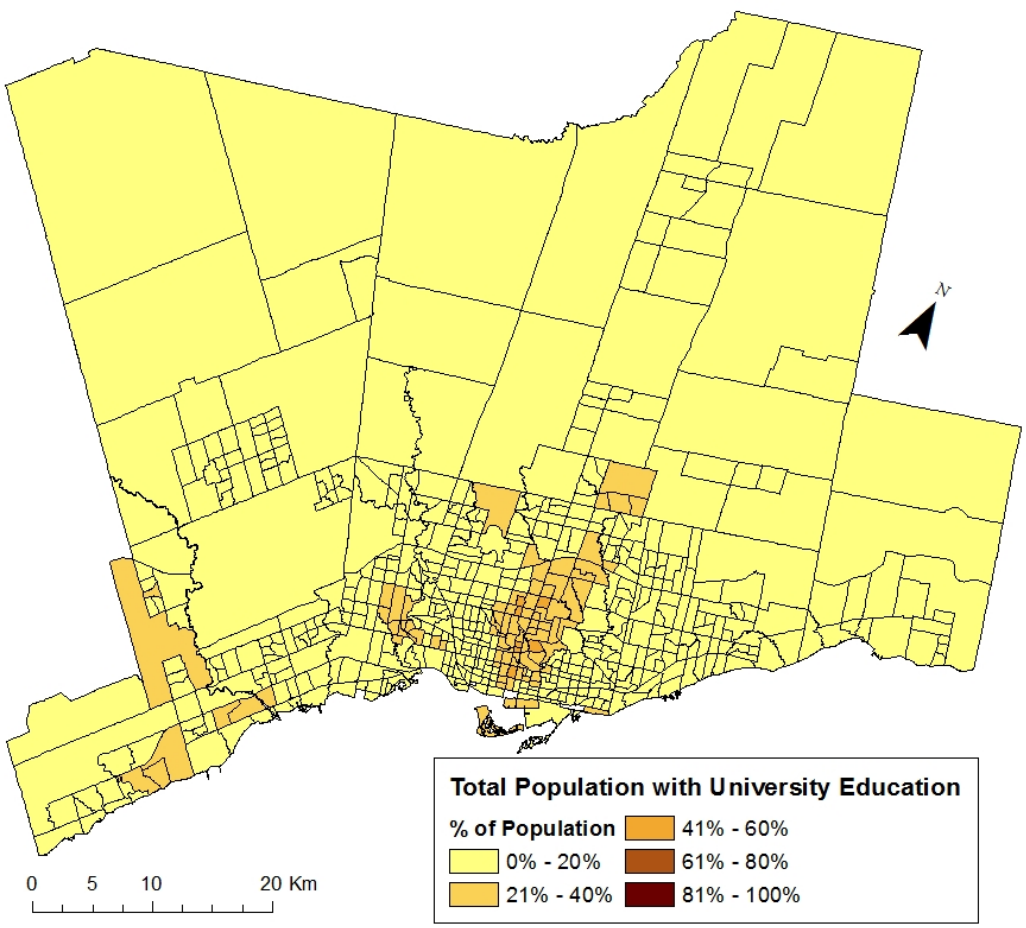

Figure 2:University Education by Census Tract across the Toronto CMA, 2016

It is always advisable to use standardized boundaries if they are available, since there is the possibility that any changes over time may be due in part to the shifting boundaries rather than actual changes over the same area. However, in many cases, data aggregated to standardized boundaries may not be available. The use of data-smoothing coupled with 2-D graphing techniques can minimize the potential error associated with changing spatial unit boundaries.

Figures 1, 2, and 3 demonstrate the advantages of graphing spatial time-series data compared with mapping it. Figures 1 and 2 illustrate university-educated individuals as proportions of population by census tract across the Toronto Census Metropolitan Area (CMA) in 1981 and 2016, respectively. Both Figures 1 and 2 illustrate similar patterns of Downtown Toronto having the highest concentrations of university-educated individuals, which generally decreases heading out into the exurbs. To be clear, it is evident that the population within the CMA became increasingly educated across all areas of the region from 1981 to 2016. However, only Figure 3 is able to clearly illustrate the relative changes in the proportion of university-educated population across different zones of the CMA within the 35-year time span. By 2016, the Toronto CMA experienced a slight increase in the concentration of university-educated individuals across the inner and outer-suburbs, from approximately 18 to 60km away from the CBD. This change is less clear when only Figures 1 and 2 are analyzed.

Both Figures 1 and 2 are useful for analyzing localized pockets of university-educated individuals, but are less adequate for summarizing changes occurring at a broader scale. The increasing number of census tracts, their changing boundaries, and the complex reality of urban geographical phenomena make visual summarizations of metropolitan-wide changes challenging to do with thematic maps. Figure 3 demonstrates how census data from Figures 1 and 2 can be graphically summarized using the LOESS method. Census tracts are plotted according to their distance from the CMA's CBD. Through scatterplot smoothing via LOESS, the increasing prosperity of the inner-city is more easily juxtaposed against the declining inner-suburbs. As well, households with higher incomes are shown to be increasingly at the suburban fringes of the CMA.

Figure 3: University Education by Census Tract vs. Distance from CBD within the Toronto CMA, 1981 and 2016

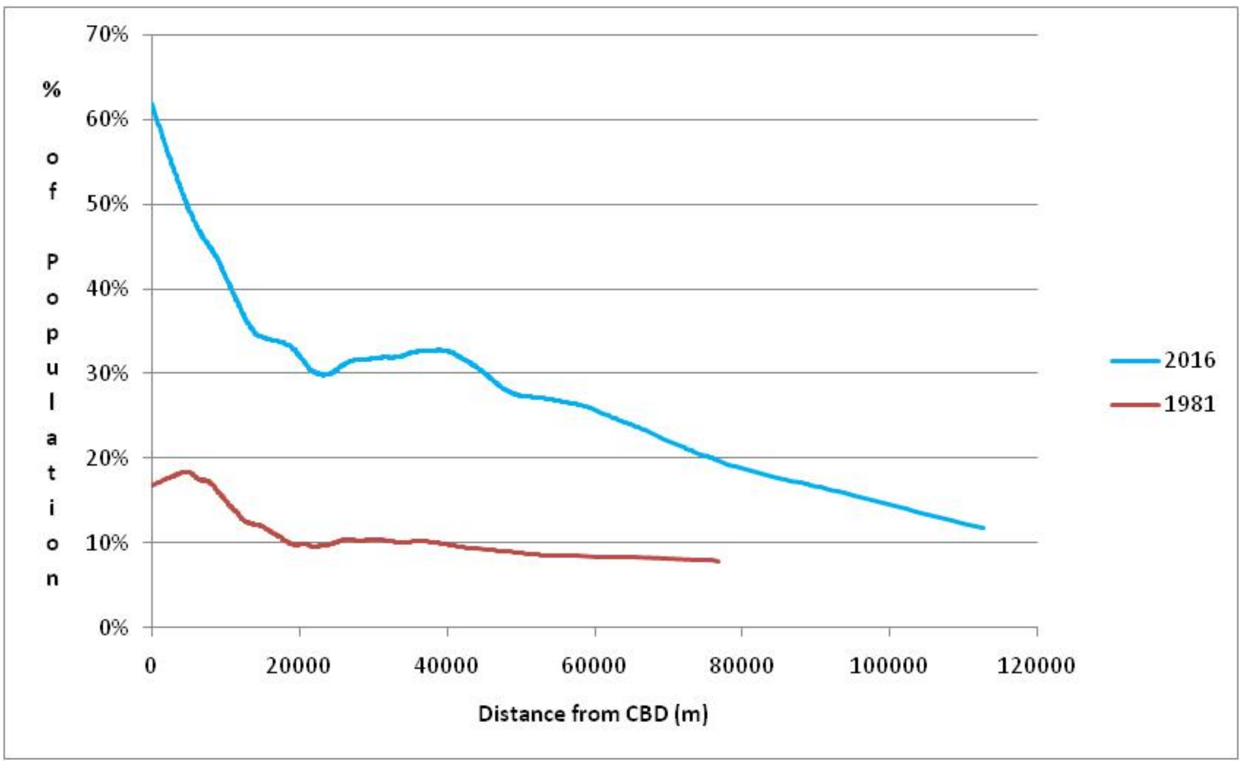

Note: the total number of university-educated individuals represented as a proportion of the total population from each census tract within the Toronto CMA was plotted on a scatterplot against their respective tract's distance from the CBD. The scatterplot smoothing method, LOESS, was then applied to the scatterplots of each year shown in the chart. The smoothing parameters (*a*) for both curves were 0.2.

3. Data Smoothing: Comparing the Alternatives

Scatterplots display bivariate or multivariate data, enabling visual assessments of relationships between variables of interest (Jacoby 2000). Visually assessing relationships however, is difficult in practice due to the presence of noisy data values, sparse data points, and weak correlations. Dealing with the previous problems involves fitting a smooth curve to the scatterplot, with two general strategies for fitting smooth curves: parametric and non-parametric fitting. Parametric fitting involves the data analyst first specifying the structure of a relationship (i.e., linear, 2nd order polynomial, logarithmic, etc.), while non-parametric fitting requires no, or very few, assumptions to be made about a relationship's structure.

For the former, specifying the "correct" structure of the relationship is almost always unknown at the beginning of an analysis, potentially resulting in a curve misrepresenting the data structure. Non-parametric fitting methods do not share that problem as they require no prior specifications – instead, a curve is calculated to pass through the densest areas of the scatterplot without constraint. Scatterplot smoothing is lauded as an essential tool for exploratory data analysis because of its simplicity and ease of interpretation towards the trends and patterns revealed (Hardle and Marron 1995). For research in the social sciences, non-parametric modelling methods are useful for summarizing relationships between social, economic, and political indices without prior assumptions about their structure.

Beck and Jackman (1998) note that while hypotheses derived from the interactions between certain variables or indices suggest the kinds of relationships that may exist between them (if any), along with their direction, there is generally little detail about their functional form. In these instances, social scientists may rely on prior assumptions and impose functional forms onto the data – for example, Beck and Jackman (1998) lament the tendency of political scientists to default to parametric linear regressions whether or not social or political theories suggest a global linear relationship between a particular set of variables. Thus, social scientists may gloss over or even exclude the possibility that there may be local variations within the functional form of a relationship (Beck and Jackman 1998). This includes local variations within metropolitan regions where variables typically show clear spatial concentrations, odd-shaped trajectories, and otherwise irregular or non-linear patterns.

To date, the application of non-parametric smoothing methods to data arranged in scatterplots has mostly not involved spatial data, but instead time-series data. The idea has been to use non-parametric smoothing to eliminate the noise in the rapid movements of a particular variable over time in order to establish the underlying secular trend. One example includes the use of the exponential smoothing method to forecast stock market volatility in fifteen stock markets (Balaban, Bayer, and Faff 2006). Another example is from Jacoby (2000), involving his demonstration of LOESS for summarizing the relationship between high-school graduation rates and voter turnouts across states in the 1992 United States Presidential Election. Non-parametric smoothing allows for the underlying secular trend in stock prices or the relationship between education and voter turnout to be graphed without bias as to any expected speed or shape to the trend (as would occur under parametric methods).

I have identified two alternative non-parametric smoothing methods commonly used for such purposes: exponential smoothing and LOESS. However, there are other non-parametric smoothers, including but not limited to weighted moving averages,[1] kernel smoothers, splines, and their variations,[2] which have been applied in various geographical settings. For example, kernel smoothing has been utilized to interpolate violent and vehicular crime using point data across selected municipalities within the United Kingdom (Martin and Ralphs 2014).

Hutchinson (1995) used thin plate spline smoothing to interpolate mean rainfall across geographic space between selected Australian climate stations, while Berke (2004) added to the health geography literature by exploring how estimates of disease occurrence or risk of disease from a regional database may be interpolated onto a continuous surface via kriging. While powerful and appropriate within their geographical contexts, the other smoothing methods previously mentioned will not be compared with LOESS in this paper. Those methods are more appropriate for settings in which data points are compared against each other in geographic space, rather than against a single point of origin. Thus, this paper will focus on exponential smoothing and LOESS. Below, I introduce each method, and discuss their strengths and limitations.

---

1 Before using LOESS as the primary data smoothing method, we tested weighted moving average for its suitability in our research. However, we found the results to be unsatisfactory as the weighting introduced a considerable amount of “delay” in the graphs that shifted census tracts further away from the CBD, thus misrepresenting the socio-spatial structure of cities. Although exponential smoothing also introduces a “delay” in the graphs, it is quite minor compared to the “delay” from weighting moving average.

2 See Rodriguez (2001) for more detail about these other non-parametric smoothers.

4. Exponential Smoothing: Definition, Usage, Advantages, and Disadvantages

In this paper, I will compare and contrast exponential smoothing against LOESS to demonstrate that the former data smoothing method is the less preferred method for visually representing geographical phenomena. Within this section, I will outline the purpose and mechanics of exponential smoothing along with its advantages and disadvantages. I begin by defining exponential smoothing, which is a data smoothing technique traditionally used to smooth time-series data. It is one of the most popular smoothing and forecasting methods, and is widely used across many undergraduate and graduate business programs (Ravinder 2013).The term “exponential smoothing” encompasses multiple methods which all fit different purposes across business, economics, and finance, including functions to fit linear, exponential, damped, constant, and seasonal trends[3] (Gardner 1985).Its popularity is due to its relative simplicity in its model formulations for short-term forecasting, and only minimal data storage and computational effort are required (Gardner 1985).

The single exponential smoothing method is a non-parametric method suitable for summarizing and/or forecasting time-series trends which have no clear seasonal or cyclical patterns. All exponential smoothing methods start with an initial value (for instance, at the beginning of a time-series), and modify that value incrementally in accordance with nearby values by assigning exponentially decreasing weights to data observations that are older (or further away) in producing a forecast for the next value (e.g. the next time period, or next place, as the case may be), such that more recent observations (or observations that are closer in space, as in our purposes) are given more weight in the forecasting process.

For single exponential smoothing, the smoothing constant, a, controls the rate at which the weights decrease, and ranges from 0 to 1.The closer a is to 0, the more weight is given to observations further away or in the distant past. That is, with lower values of a, more of the entire dataset is taken into account when smoothing out the local average. The closer a is to 1, the more weight is given to the most recent observations, which means that data points most

3 Twelve methods are used to fit different short-term forecasting models. For their functions and form, see Gardner (1985).

close together (in time, or space) are taken into account in forecasting the local average around each point. The formula for the single exponential smoothing method is as follows:

$$F_{t+1+} = aY_t + (1 - a)\widehat{Y}_t$$

where Ft+1 is the forecasted value for the immediate upcoming period, a is the smoothing constant, Y_t is the actual explanatory variable value for the current period, and $\widehat{Y}t$ is the predicted explanatory variable value for the current period. I have produced a hypothetical of single exponential smoothing and use of the smoothing parameter *a*, shown in Figure 4, which graphs the hypothetical price per share (USD) in 2020 of a hypothetical company X.

Notably, as illustrated in Figure 4, the exponentially smoothed curve "lags" in its representation of the time-series trend relative to the actual data points. The method misrepresents the stock price of Company X as peaking in September, when in actuality the company's stock price peaked in July. This "lag" is found in all applications of exponential smoothing. Within the context of spatial time-series data, the noted "lag" becomes less noticeable across vast geographical distances.

Figure 4: Hypothetical Exponential Smoothing as Applied to Monthly Share Prices for Company X*

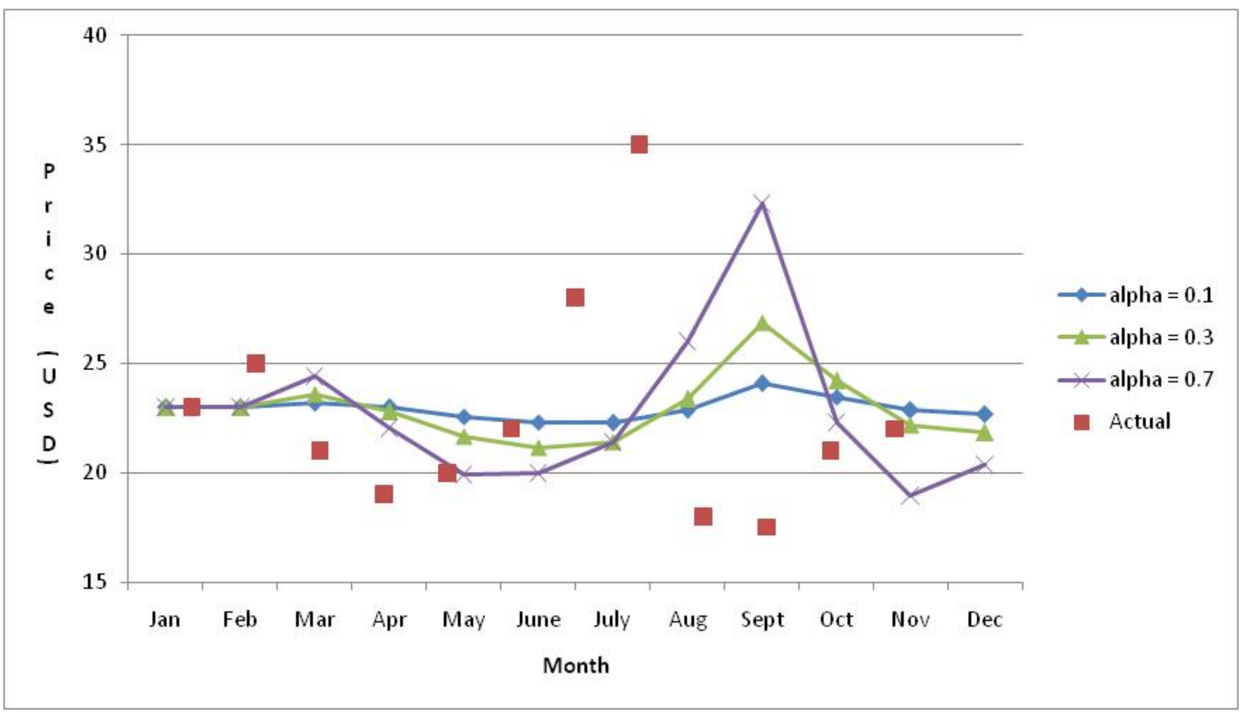

Note: The curves in Figure 1 represent different single exponential smoothing solutions, each using a different smoothing parameter (*a* = 0.1, 0.3, 0.7). Under the exponential smoothing method, *a* values closer to 0 produce curves which are smoother but flatter, while values closer to 1 produce curves which follows the original data points more closely.
*Company X and its hypothetical share prices for 2020 are completely fictional, and any resemblance in performance to a non-fictional corporation is completely coincidental and unintended.

The smoothing constant, *a*, can be optimized to ensure forecasts are as accurate as possible. Summary error metrics such as Mean Absolute Deviation (MAD), Mean Squared Error (MSE), or Mean Absolute Percent Error (MAPE) are used to optimize forecasting accuracy (Ravinder 2013).The optimal value of *a* is one which minimizes MAD, MSE, or MAPE. Every exponentially smoothed curve will have its own optimal a to solve for. The extreme values (0 or 1) for a are of concern if and when they appear, discussed below. Separate formulae that solve for an optimal value of *a* thus become necessary to minimize MAD, MSE, or MAPE, and contemporary software has increasingly included solver functionality to allow for this.[4] Figure 5 from Ravinder (2013) summarizes how MAD or any other error summary statistic is used to determine the optimal *a*.

Figure 5: The Relationship Between Mean Absolute Deviation (MAD) and Different *a* Values

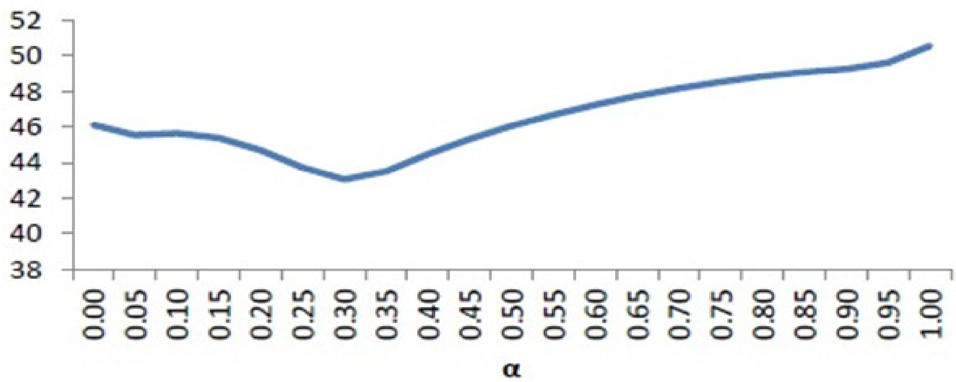

Source: Ravinder (2013)
Note: the optimal *a* is the one where MAD or any other error summarization statistic is minimized, which in this example would be at the point where *a* = 0.30.

The advantages of single exponential smoothing are clear. The simple formula and limited computing power needed makes it an attractive smoothing and forecasting method for time-series analyses. Yet, there are several disadvantages of this method that reduce its usefulness for data arranged in certain ways, and that make it less applicable to urban geographic data. The first issue was pointed out by Ravinder (2013) when he found many that many statistical problems had an optimal *a* of either 0 or 1. These extreme values are undesirable because they often do not adequately summarize the data, producing either curves that are not smooth (*a* of 1) or that are just straight curves with little relevance to the data (*a* of 0). Ravinder (2013) hypothesizes that these instances where the optimal a was 0 or 1 are a result of two potential data limitations: (1) insufficient time periods, which do not allow sufficient adjustment for the exponential smoothing to show results, and (2) having an initial data point that is very close to

4 The Solver function of Microsoft Excel (2010 and beyond), because of its speed and the possibility for automation, was used to solve for the optimal *a* when attempting to produce smooth curves through distance-based scatterplots via exponential smoothing for spatial time-series analyses. The non-linear optimization algorithm of versions of Excel from 2010 and beyond were found by Ravinder (2013) to have correctly determined the optimal *a* in 86 percent of the problems he reviewed related to single exponential smoothing. The Solver function in earlier versions of Microsoft Excel were found by Ravinder (2013) and other scholars to be unreliable in finding the optimal *a*. These versions of Solver assumed linear relationships between *a* and error values, thus the optimal *a* would always be at one of the extremes. They did not consider the possibility of non-linear relationships, such as the one in Figure 2.

the average of the rest of the data. He argues that insufficient time periods make it difficult to see a trend, let alone changes in the trend.

A value of 1 may arise because the exponential smoothing was unable to detect any changes in the trend due to the smoothness of the data points. A value of 0 may arise if, in addition, the initial data point observation is very close to the average of the data, which then does not induce the method to predict any changes in the forecast value over time. In such cases, MSE or MAD is minimized if each subsequent forecast is kept at or close to the average of the data, leading to minimal changes or no changes to the initial forecast. The initial forecast must be different enough from the data points that follow it, and there must be enough observations for the trend to change over the course of the scatterplot (which typically means over time). Thus, there should be a lower likelihood of the optimal a being 0 or 1 when there is sufficient data.

In working with the Canadian census data, we found evidence corroborating Ravinder's hypothesis when applied to spatial data. Smaller metropolitan regions posed situations where the optimal solution for $a = 0$, with a result of no changes in trends despite obvious and clear secular trends in the scatterplots. In working with smaller datasets, this disadvantage to the exponential smoothing method becomes especially evident and disruptive.

The second disadvantage of single exponential smoothing is its greater sensitivity to sharper changes in values which often produce relatively unsmooth curves (at least, without further smoothing parameters, but this questions the point of using the exponential smoothing if additional smoothing methods are required). At the same time, if there is already a smooth secular trend that does not involve sharp changes in underlying values, combined with having a dataset with fewer observations, the exponential smoothing method may not sufficiently track the secular trend (which often results in $a = 0$ when optimized).

The third disadvantage is that each problem or scatterplot will have its own optimal a value that maximizes the efficiency of the exponential smoothing parameter. This means effectively that each exponentially smoothed curve will have used a different method for tracking the secular path through a series of observations (those with higher a values will have smoothed curves that adapt to changes in the underlying data more quickly than those curves using lower a values), calling into question the comparability of the results. The only solution to this is to use the same a value for every smoothed curve one wants to compare (for example, one could use the average a value as solved for a series of different datasets). However, this is problematic as it means some of the resulting smoothed curves will actually be more "true" to the underlying data than others. Furthermore, this means some curves will be either more overfitted (or smooth) relative to others in the same situation, or underfitted (insufficiently smooth, thereby reducing the effectiveness of the exercise).

The fourth disadvantage of single exponential smoothing is that because the method fits one curve through the dataset one single observation after the other, it is not well adapted to the application of differential weights among the data points. While this is not a problem with single observations that all have equal weight (such as end-of-day share price indices), it is a problem for observations that need to be differentially weighted. When the data points are neighbourhoods, for instance, it is useful and necessary to weight each data point by its relative population, so that under-populated neighbourhoods do not disproportionately affect the path of the smoothed curve through the scatterplot.

Figure 6: LOESS vs. Exponential Smoothed Curve with Spatial Data: University Education by Census Tract, Regina CMA 2001

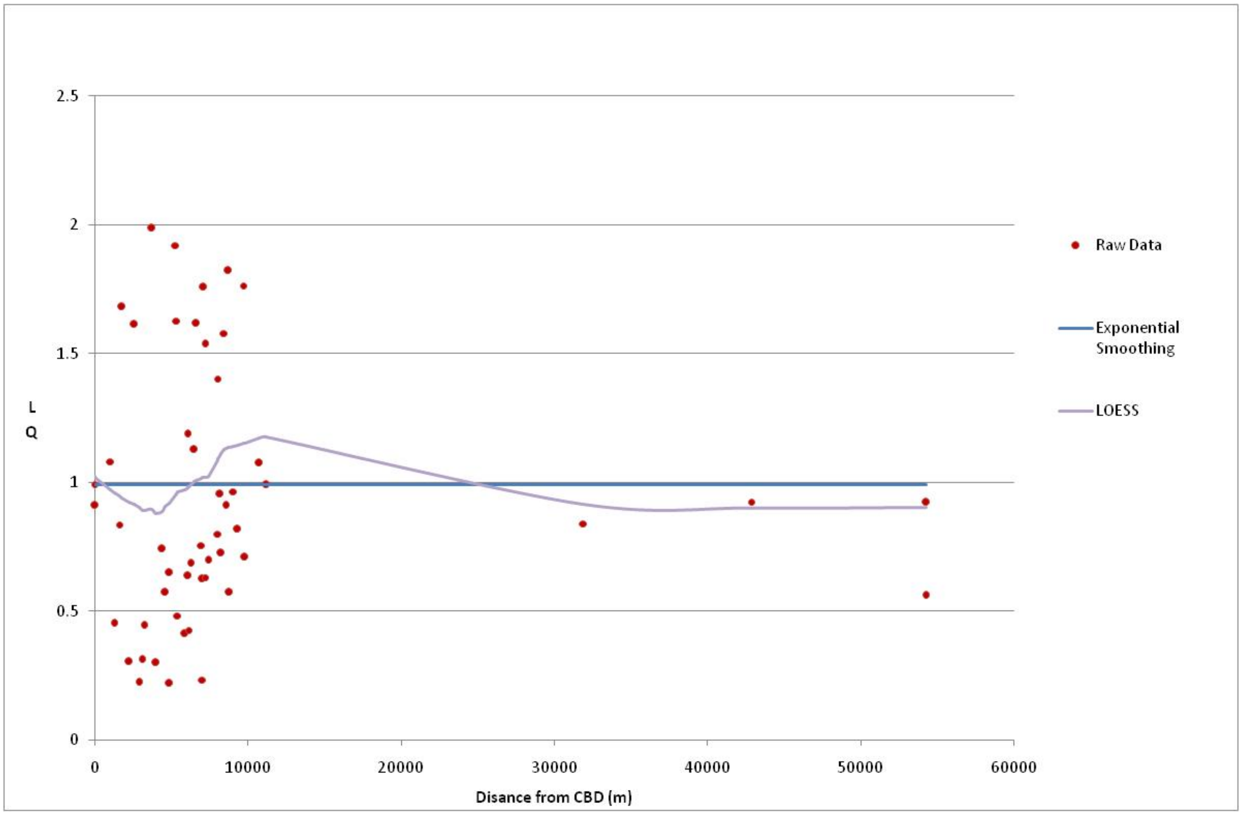

Note: Raw data values represent the spatial distribution of the Regina CMA's population over 15 years of age with a university degree, certificate, or diploma in 2001.The smoothing parameter (*a)* used for the exponentially smoothed and LOESS curves is 0 and 0.5, respectively.

Figure 6 demonstrates clearly three of the four weaknesses of exponential smoothing noted above, in a graph depicting its fitness against raw data for educational attainment by census tract in the Regina Census Metropolitan Area (CMA) in Saskatchewan, with census tracts arranged by distance to the central business district (CBD). Figure 6 demonstrates the effects of a small dataset, an ambiguous yet relatively smooth secular trend in the scatterplot, and the detriment of lacking weighted data points for the exponentially smoothed curve. The exponentially smoothed curve does not sufficiently track the underlying trend of the data as there are no sharp changes across the spatial data, with the resulting optimized smoothing parameter being 0.Without weighting the data points by their relative populations, there are no "pulls" on the exponentially smoothed curve. The result of all these factors is a straight line through the data points. However, the trends, while not immediately obvious, in actuality do not amount to a straight line. There is a slow increase in the average concentration of university-educated individuals as one moves outwards from the CBD, which then slightly declines as one moves into the exurbs. The correct functional form of this relationship should somewhat approximate a 3rd order polynomial curve. Any alternative non-parametric smoothing method must be more responsive and applicable to relatively smaller datasets such as these.

5. Locally Weighted Scatterplot Smoothing (LOESS): Definition, Usage, Advantages, and Disadvantages

The most common alternative to exponential smoothing is locally weighted scatterplot smoothing, or LOESS. LOESS is a powerful, non-parametric modeling method for curve fitting. It is currently the most popular non-parametric smoothing technique (Harrell 2015; Jacoby 2000).LOESS is often used for exploratory analysis, in particular for fitting smooth curves to scatterplots for the purposes of data visualization and interpretation, and is considered among the best for allowing "the data to speak for themselves." (Jacoby 2000). Like the exponential smoothing method, the LOESS procedure does not require any prior specifications about the data structure to be made. Importantly for geographers, I have found that fitted LOESS curves can reveal local variations critical to understanding inter-zonal trajectories among variables across metropolitan space that would otherwise not be easily detected.

LOESS operates differently from the exponential smoothing technique. Instead of incrementally modifying the previous forecasted value using the information from current observations, as the exponential smoothing method does, the LOESS method estimates new forecasted values without bias as to what the previous forecast had been. This eliminates entirely the problem noted above for exponential smoothing in which an initial value close to the dataset average could produce no trend, while avoiding the other problems with the exponential smoothing method that were previously discussed.

There are some limitations, however. While the algorithm for calculating LOESS is relatively simple compared with other non-parametric regressions, it is a computationally intensive method (Engineering Statistics Handbook, n.d.).However, this drawback is not particularly problematic unless the datasets of interest are extremely large. Another drawback of LOESS is that it requires a minimum size of data points, with fairly densely sampled datasets, in order to produce estimates (Engineering Statistics Handbook, n.d.), which places a floor on the size of dataset one can use with LOESS. However, the LOESS method can handle datasets that are smaller than those required for full functioning of the exponential smoothing method. A final drawback is that the LOESS method does not produce a regression function that is easily represented by a mathematical formula, making it difficult to transfer or communicate the results from one solution to others. Those who wish to view the regression function will have to have the dataset, user specifications, and software used to calculate LOESS (Engineering Statistics Handbook, n.d.). However, each of these drawbacks is relatively minor compared to the benefits that LOESS provides for analyzing spatial time-series data.

LOESS performs curve fitting by combining the simplicity of least-squares regression with the flexibility of nonlinear regression (Engineering Statistics Handbook, n.d.).Simple regressions are fitted to localized subsets of the data to build up a function describing patterns of localized variation exhibited across the scatterplot, point by point.[5] LOESS begins by selecting a series of *m* locations or evaluation points, v_j, where *j* ranges from 1 to *m*. The evaluation points are equally spaced across the range of X,[6] but located relatively close to one another such that the locally fitted curves connect to form an overarching smooth LOESS curve (Jacoby 2000).At each point a linear or low-degree polynomial (usually 2nd degree) is fit to a subset of the data, although the user may also specify the regression to be performed.

Simple regressions such as linear or 2nd-order polynomials are usually fine for approximating any local function and easily fit data within small subsets (Engineering Statistics Handbook, n.d.).A rule-of-thumb is that should the scatterplot adhere to a monotonic pattern, the model should be set to linear, while a non-monotonic pattern merits a quadratic model (Jacoby 2000).The user defines how many data points (the smoothing parameter, *a*) are included in a subset.[7]The smoothing parameter is the proportion of all data points in the dataset to be included in each subset. Larger values of *a* will produce smoother LOESS curves which fluctuate the least in response to outliers or noise, while smaller values will produce curves conforming more closely to the data, which may not be as desirable if they start to capture random error or noise in the data.

The data points included in each local regression are inversely weighted according to their distance from their evaluation point, v_j. This weighting is grounded in the concepts of spatial autocorrelation and spatial dependence. According to Jacoby (2000), LOESS can be conceptualized as a “vertical sliding window” moving across the X-axis of the scatterplot, stopping at each evaluation point and fitting simple regression lines across the subsets, whose widths are defined by the smoothing parameter, a. The sliding window means that each local regression only includes data points captured within the window, allowing the estimated slopes of the regressions to follow the contours of the data, giving LOESS its characteristic flexibility (Jacoby 2000). Because of this, the LOESS smoothing method is flexible and unconstrained by older data points or those further away.

---

5 For an in-depth analysis behind the mathematical steps involved in fitting a LOESS curve, see pages 608–12 in Jacoby (2000).

6 Jacoby (2000) states that the exact number of evaluation points is “relatively unimportant, so long as there are enough of them to provide sufficient detail about the variability in the conditional distribution of the Y variable.” The value of m is usually determined by the software used to calculate LOESS.

7 An additional robustness step in LOESS may be included, but is optional for calculating LOESS. Its purpose is to reduce LOESS sensitivity to outliers by down weighting observations with large residuals; however, extreme cases can still overcome the method (Engineering Statistics Handbook, n.d.) See Jacoby (2000), pages 587–90, for more detail regarding the effects of the robustness steps on LOESS calculations and results, and how robustness is calculated.

6. Applying the Alternatives: A Canadian Case Study

As mentioned above, the aim here is not only to ascertain the best method for use in time-series analysis, but to use non-parametric smoothing for comparing spatial data series, including changes in the patterning of spatial data over time. To this end, we have applied and compared these two methods to spatial time-series data for metropolitan areas, using Canadian metropolitan areas as the case studies. We are here interested in examining how non-parametric methods may be used for detecting relative changes in clusters of neighbourhoods grouped via spatial locations.

The figures and examples presented in this paper were derived from Canadian census data. Canadian metropolitan regions are defined and bounded by Statistics Canada's Census Metropolitan Areas (CMAs), inside which are smaller census enumeration units such as census tracts, the unit of analysis for the spatial time-series data we analyze herein. To capture CMA dynamics and trajectories across different points in time, census data was collected at the tract level from the 1971, 1991, and 2016 Canadian Censuses conducted by Statistics Canada. Unfortunately, Statistics Canada often changes the boundaries of their census tracts, particularly when the local population grows inside one, in which case it may be split into two or more tracts, or when population expands at the edges (in which case the outer boundaries of existing tracts may be modified).

Thus, the census tract dataset used here involves unstandardized boundaries. The LOESS method can be applied using these unstandardized boundaries to produce a properly fitted curve.As I am interested in comparing trends over time, I sought to remove as much error as possible associated with changing spatial unit boundaries. To partially overcome this limitation, values for census variables can be interpolated between census tracts using the tracts' distances from the CBD.[8] That is, rather than having discrete values for census variables demarcated by census tract boundaries, a continuous gradient of values was created by interpolation using the tracts' distances from the CBD. However, shifting census unit boundaries means a spatial shift in their centroids' distances to the CBD, which then has the potential to change the spatial distribution of the explanatory variable and introduce distortions into the analysis. Nevertheless, in many cases the only data available is arranged in spatial units like census tracts and the boundaries may change. Being able to deal with this through localized parametric techniques is one of the strengths of the LOESS method, although if boundary shifts are not accompanied by changes in where the population is located within a tract, this could potentially affect the results.[9]

Similar to a time-series graph, census variables were graphed separately for the six points in time related to each census in order to analyze changes across time. In each case, the x-axis delimits the distance from the CMA's CBD.[10]Across the CMAs of interest, census tract distances (in metres) to the CBDs were calculated by taking the distance from their centroids to the CBD's centroid. The distances were calculated for 1971, 1981, 1991, 2001, 2006, and 2016.The census variables were then graphed against distance on scatterplots by year.

Six LOESS curves were fitted according to each census year for each dependent variable. Linear functions were fitted in each local regression. After trial-and-error, a smoothing parameter of 0.2 was chosen to be used for the LOESS analyses related to examples provided herein (higher values create smoother lines, lower values bumpier lines). As neighbourhoods are the scale of interest in this project, census tracts are the main census enumeration unit for analysis. They are small, stable geographic areas, created to be proxies for neighbourhoods by Statistics Canada for the purposes of reporting census data, and which are bounded by local rivers, lakes, and major roadways and railways.

Variables included from the Canadian censuses were processed into location quotients (LQ) for ease of interpretation. LQs measure the concentration of a variable in a census tract relative to its concentration across the entire CMA.LQs range from 0 to infinity, where 1.00 indicates an identical concentration between a census tract and the CMA average, while values above or below 1.00 indicate greater or lesser concentration respectively. LQs for each census tract are derived from the following formula:

8 Although distance from CBD remains a useful metric in representing the socio-spatial structure of cities, it is limited by its monocentricity, which does not represent the multi-modal and fragmented reality of modern cities. Unlike the pre-industrial or industrial cities, modern cities contain multiple nuclei acting as centres of employment and public life.

9 If, for example, a number of census tracts were to be significantly enlarged while the spatial distributions of the populations living within the tracts remained unchanged, then a time-series curve would be produced which illustrates the spatial distribution of a census variable becoming increasingly dispersed over time, despite there being no true population movement across census tracts. This is because the centroid of a census tract shifts if the tract is enlarged. In other words, shifting census tract boundaries without a corresponding shift in spatial population distributions within tracts will result in a misrepresentation of the spatiotemporal distribution of a census variable.

10 The CBD of each CMA was determined through experience and by selecting among those census tracts within or near the CMA's downtown that had the highest concentration of the tallest office buildings.

$$LQ = (X_i / Y_i) / (X_{cma} / Y_{cma}) \quad i=1$$

where i is the census tract in question and cma is the CMA in which the tract is located. X and Y represent a very wide range of variables. X represents the sub-categories of census variables, such as university educational attainment, occupational grouping (e.g., manufacturing), etc. Meanwhile, Y represents broad census variables which act as the base population against which each sub-category in X is compared (e.g., all adults aged 15 and over, when examining university attainment, or the total labour force when examining a particular occupational group, etc.).

Prior to fitting LOESS curves to the scatterplots, all census tracts were weighted according to their relative populations. Such weighting is required so that the more populated census tracts have greater influence over the resulting LOESS fitted curves, while smaller tracts would have less influence. This is a major benefit of the LOESS method. One issue related to data smoothing of LQs in particular is that they are not bounded on the upper end, and may be positively skewed. Because of this, the observation values input into the models were capped at $LQ \leq 8$ to prevent extreme values from having disproportionate effects on the resulting curves. This step further reduced the influence of extreme values on the LOESS or exponential smoothing.

For ease of replication and speed of the procedure, LOESS was automated and performed using Microsoft Excel 2007, the coding of which was developed by Peltier Technical Services.[11] The goal for future exploratory analyses is that any data smoothing method be easily replicable, editable, automated, and quick in particular. Hardle and Marron (1995) agreed that speed is important for exploratory analyses since “too long a lag between conception of the ideas of the experimenter and their visual realization hinders the analytic process” (p. 1).

After trying a variety of solutions, a smoothing parameter of 0.2 was settled on as it best fit the datasets in which census tracts in the large Canadian CMAs were the units of analysis, while 0.5 was used for small CMAs. Furthermore, this smoothing parameter was then applied to all the CMAs in the study for each study year, in order to maintain consistency in the approach, and to make sure spatial autocorrelation was treated equally across CMAs by the method. For the purposes of spatial time-series analysis using Canadian census variables to graphically summarize metropolitan dynamics and trajectories, consistency is critical to the analysis. Having ubiquitous smoothing parameters ensures each CMA across all time periods are smoothed by the same values and that dynamics, patterns, and trajectories noted by the user are the result of underlying trends in the data, rather than of changing smoothing parameters.

With regards to the optimization method detailed by Jacoby (2000), the distance-based analysis undertaken in this research project introduced spatially autocorrelated trends into the variables, creating residuals that reflect those trends. Thus, any LOESS curve fitted to the residual scatterplots will always be trended. The spatial-autocorrelation introduced by the distance-based analysis is not an undesirable quality to be removed from the analysis, as the entire point is to detect localized spatial dynamics, patterns, and trajectories using LOESS. A

11 For complete details on the mathematical procedures coded into the LOESS Excel add-in, in addition to its VBA coding, see Peltier (2009).

consistent smoothing parameter is therefore desirable when conducting comparative analysis of spatial relationships, and when comparing patterns over different time periods, as it means that these spatial relationships are treated the same for every CMA.

Smoothing parameters of 0.2 and 0.5 were adopted for the case-study CMAs (Toronto and Regina) modelled in this paper. This was determined via trial-and-error to offer a satisfactory combination of curve smoothness and data fitting. However, the smoothing parameter under LOESS can also be subject to an optimization method such as the one outlined by Jacoby (2000) to prevent data overfitting. Jacoby (2000) elaborated upon a method which saw the user plot the residuals from LOESS on a scatterplot, and fitting a LOESS curve (with the same smoothing parameter as the original) onto the residuals. The resulting LOESS curve should be relatively straight and horizontal, indicating that most of the functional dependencies between the X and Y variables have been picked up by the original LOESS curve. However, due to the importance of using ubiquitous smoothing parameters and the introduction of spatial-autocorrelation into the census variables, this optimization is less applicable to my purposes in tracking census tract variables across geographic space.

7. LOESS vs. Single Exponential Smoothing

The differences in the effects of using each smoothing method can be ascertained when applied to datasets for large CMAs. Here, the Toronto CMA, which is Canada's largest metropolitan region with the largest population, acts as the main case study. This CMA is large enough that the exponential smoothing method delivers an optimal a solution that is not 0 or 1, thus avoiding one of the key problems of the exponential smoothing method that arises when smaller datasets are analyzed. Below, LOESS is compared against exponential smoothing in this section using an education variable: proportion of population aged 15+ that is university-educated. First, curves fitted using LOESS are contrasted against those fitted using exponential smoothing, using the education variable from the 2016 Census of Canada for the Toronto CMA (Figure 7).

As can be seen from Figure 7, which shows university education on the y-axis, and distance from the CBD on the x-axis, the LOESS fitted curve is much smoother than the curve produced by single exponential smoothing. The latter remains overfitted, and still communicates some of the noise inherent in the rapid shifts in the average values over space. Furthermore, because the exponential smoothing method fits its curve incrementally with each new observation (in this case, starting from the CBD and incrementally changing the parameter with each new observation to the right, until it reaches the urban-rural fringe of the CMA), it moves much more slowly when there are fewer observations such as found on the right-hand side of the graph (in the exurbs).

As one moves from the CBD to the suburban fringe, the population densities, and the number of census tracts, diminish. The exponential smoothing method is shown to be unable to adjust fast enough, even in Canada's most populous CMA, and overestimates the actual education levels of those living in the most distant census tracts. The LOESS method, on the other hand, is not dependent on incremental changes from one observation to the next, but instead on separately fitted local regressions, which means it can adjust instantaneously. Furthermore, the LOESS method is adept at handling the smaller number of observations found here near the edge of the metropolitan region. And because the LOESS method allows for population weighting of each census tract, it is more sensitive to where the people are actually located. The LOESS method in all cases performs better here than the exponential smoothing method.

Figure 7: LOESS vs. Exponential Smoothed Curves with Spatial Data: University Education by Census Tract, Toronto CMA 2016

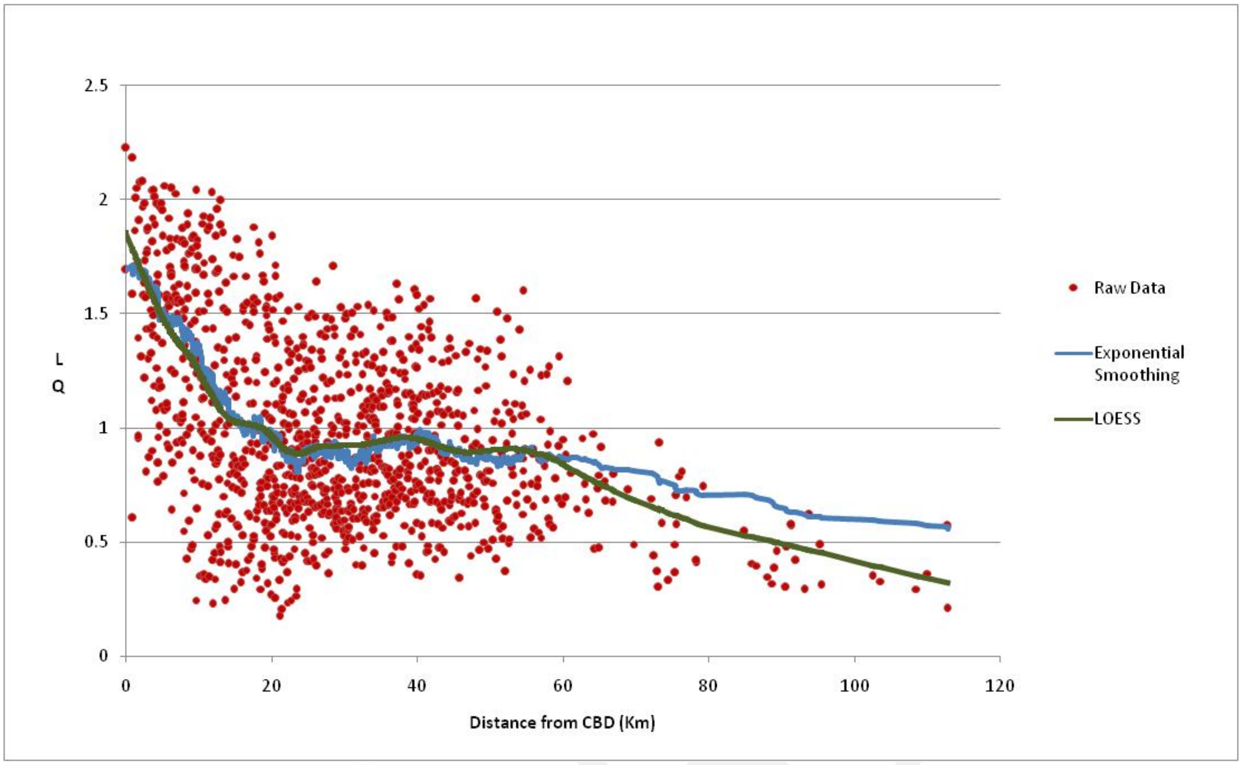

Note: Raw data values represent the spatial distribution of the Toronto CMA's population over 15 years of age with a university degree, certificate, or diploma in 2016.The smoothing parameters (*a)* used for the LOESS and exponentially smoothed curves are 0.2 and 0.031 respectively.

Not only are these smoothing methods applicable to spatial data like those in Figure 7, but they can be used to compare spatial trends over time. In Figures 8 and 9, the same census variable (LQ for the proportion of the population aged 15+ that has a university degree) is graphed for the period from 1971 to 2016, using LOESS and exponential smoothing.

Figure 8: LOESS Curves with Spatial Data: University Education by Census Tract, Toronto CMA 1971 to 2016

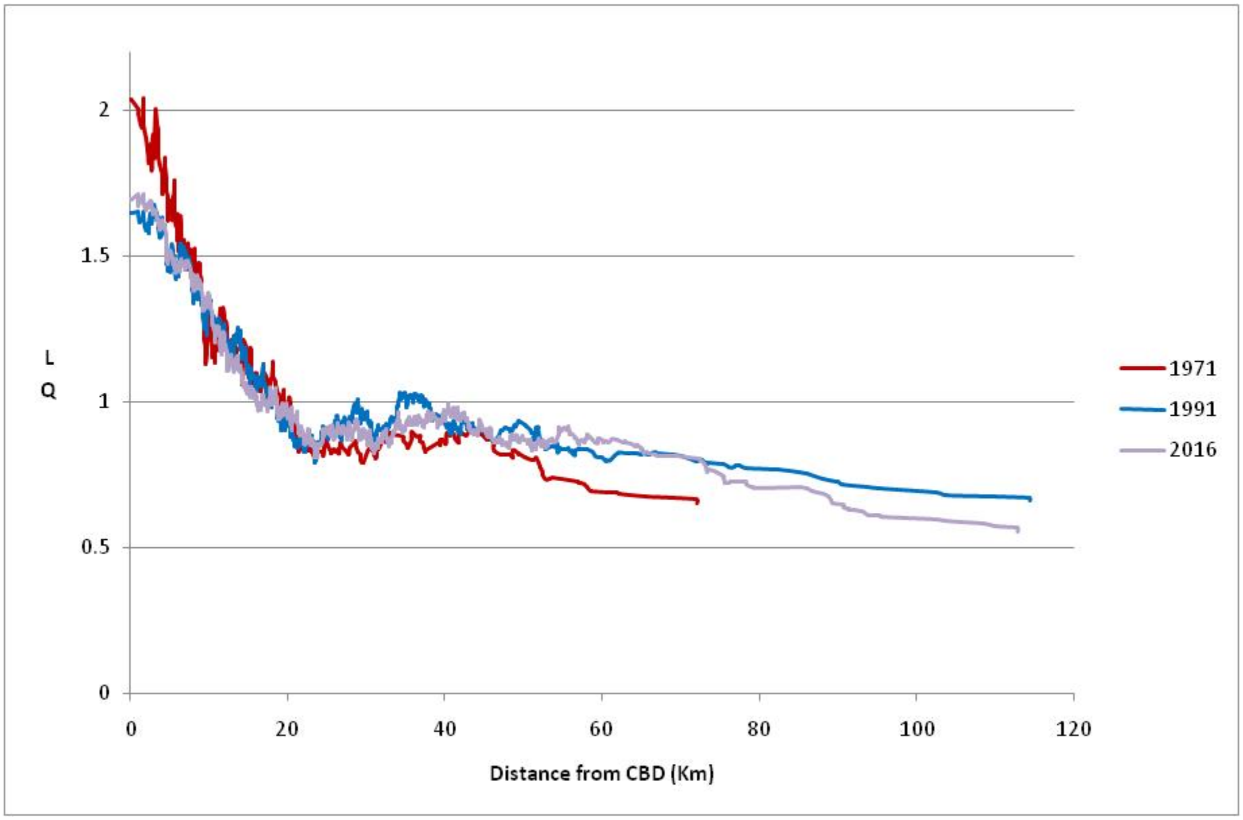

Note: The smoothing parameter (*a*) used for the LOESS curves is 0.2.

In Figure 8, the LOESS method produces a series of smooth curves that span the distance from the CBD to the urban fringe as it existed in each census period under study. There is clear differentiation amongst the fitted lines, allowing the viewer to clearly see the shift occurring in the inner-city tracts near the CBD (which went from having relatively lower educational attainment in 1971 to increasingly higher educational attainment through the 1990s and 2000s). Furthermore, shifts occurring in neighbourhoods further away from the CBD are also clear, including the initially lower levels of educational attainment in the outer suburbs and exurbs, the increase in educational attainment through the 1990s there, and then a slight decline in educational attainment in outer areas (compared to the CMA average) through the 2000s.

Figure 9: Exponentially Smoothed Curves with Spatial Data: University Education by Census Tract, Toronto CMA 1971 to 2016

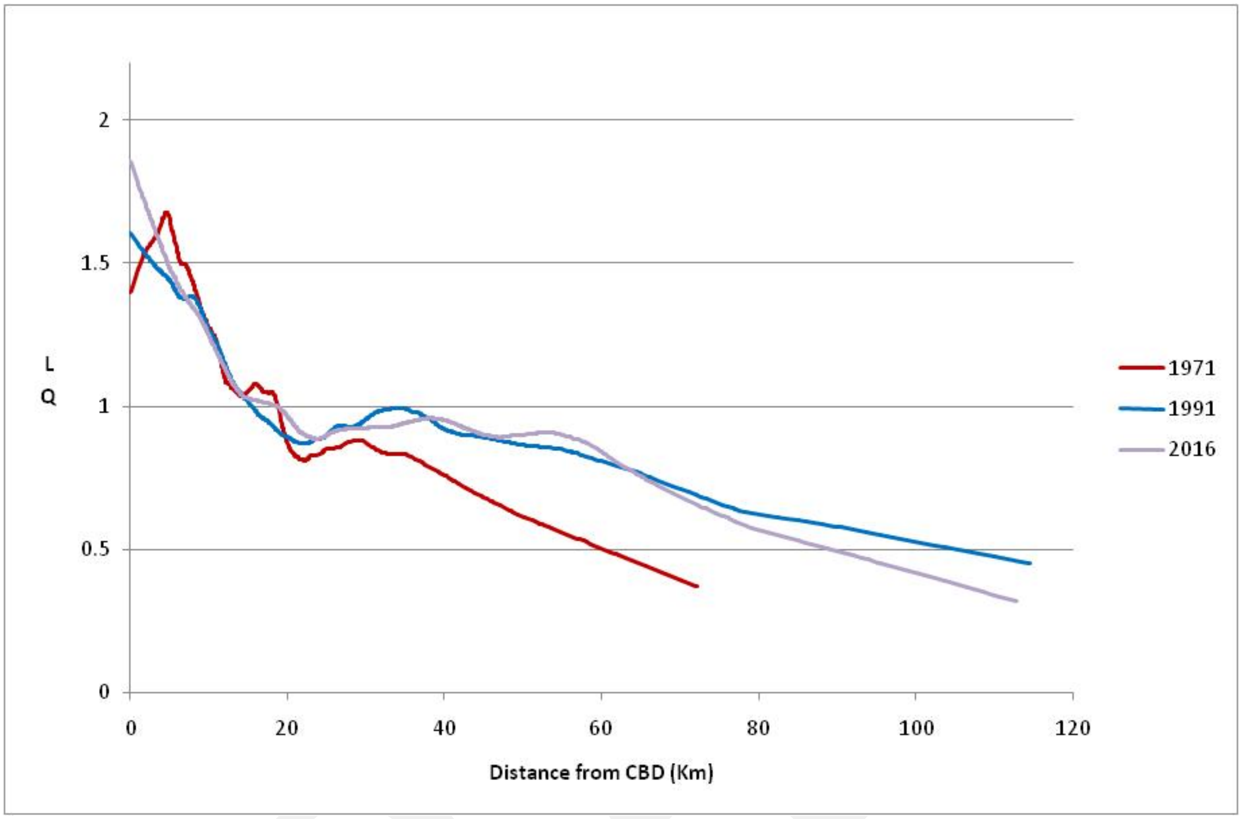

Note: The smoothing parameters (*a)* used for the exponentially smoothed curves are as follows:1971, 0.018; 1991, 0.030; 2016, 0.031.

The exponentially smoothed curves, despite having their smoothing parameters optimized, remain overfitted across all five time periods and communicate much of the noise inherent in rapid shifts in average values over space, particularly where the bulk of the population and census tracts are most concentrated. In fact, the curves are overfitted to the point at which it is difficult to make out the differences across the time periods, since they overlap so often. The overfitted curves muddle the trends of their underlying data. Furthermore, the decline in education attainment as one moves outward is minimized, and the shifts after 1991 are far more muted. In contrast, the LOESS curves are much smoother and not overfitted, and consequently overlap much less than the exponentially smoothed curves.

The result is that one can incrementally analyze, with more clarity, the changes in the underlying variables, in this case, local concentration or levels of university-educated individuals from 1971 to 2016 across census tracts in the Toronto CMA. The concentration of university-educated individuals has increased across much of the suburbs from 1971 to 2016, while also concentrating within tracts closer to the CBD. This trend is very much obscured when analyzed from the exponentially smoothed curves. The inability to utilize population weighting in exponential smoothing is very likely to also be contributing to the differences between the two methods. Altogether, Figures 7, 8, and 9 demonstrate the superior performance of the LOESS method over the exponential smoothing method.

Overall, the LOESS and exponentially smoothed curves followed similar paths through their datasets, but with significant differences between the methods. A noticeable difference was the relative smoothness of each curve: despite having optimized smoothing parameters, the exponentially smoothed curves were significantly less smooth than the LOESS curves – these curves were overfitted. We should note that this remains so, even when (or if) a smoothing parameter a closer to 0 is used instead of the optimized parameter (and even then this produces a line that is otherwise too horizontal than the underlying data).

Comparing LOESS and exponential smoothing serves as cross-validation in that the same variables were analyzed using the same dataset. The fact that the LOESS curves revealed similar trends to those of the exponentially smoothed curves serves as validation of both LOESS for non-parametric smoothing of spatial time-series data, and the use of ubiquitous smoothing parameters in analyzing metropolitan dynamics and trajectories over space and time.[12] The use of both methods offers possibilities for further cross-validation of their results in potential future tests.

12 Exponential smoothing methods have parameters that can be statistically optimized to best represent the functional form and trends of a dataset. Thus, using optimized exponentially smoothed curves should hold more certainty than LOESS in that the curves fitted to the data should accurately follow the data's trends and structure. Therefore, this comparison of LOESS to exponential smoothing served as a test of its curve fitting capabilities, in addition to demonstrating the weaknesses of exponential smoothing.

8. Conclusion

LOESS is a powerful non-parametric smoother that is useful for modelling the trends and structure of bivariate scatterplot data. As a non-parametric smoother, LOESS is flexible in modelling the nature of relationships between variables. The localized regressions across subsets of data points employed by LOESS reduce the sensitivity of the method to extreme values, and are one reason for the method's flexibility. At the same time, its flexibility is also a potential weakness in that it cannot provide a function for the data via a simple equation, meaning that other users would require the dataset, user-specifications, and statistical program used to replicate any results, and users must make partially arbitrary decisions about the fitting parameters in the absence of well-defined optimization methods.

For the purposes of graphically summarizing the relationships between social, political, and economic indices for spatial time-series analyses, LOESS is shown to be a promising smoothing method. The traditional method of data visualization for spatial analysis within the social sciences has been thematic mapping. For certain time-series analyses, mapping is a powerful method for visualizing changes across space and time – for example, to show the urban growth of a metropolitan region or to demonstrate the shrinking area of a large body of water. However, mapping could fall short when it comes to the analyzing and comparing various social, economic, and political changes across metropolitan regions over time. 2-D time-series graphs can be useful for summarizing bivariate relationships across time, and as demonstrated herein can be adapted to spatial data series of various sorts.2-D non-parametric modelling techniques allow the user to visualize and approximate metropolitan dynamics, trajectories, and changes across time more accurately.

LOESS was compared with another non-parametric smoothing method: exponential smoothing in its most common variant (the single exponential smoothing). All exponential smoothing methods assign exponentially decreasing weights of importance to its older (more distant) observations, giving recent observations more weight. Comparisons between fitted curves produced by the two methods decidedly revealed LOESS to be the superior non-parametric smoothing technique. LOESS curves were smoother than exponentially smoothed ones, and were just as accurate, if not more in terms of summarizing the underlying data trends and structure. However, this comparison was also an opportunity to perform cross-validation of LOESS' data fitting capabilities. The fact that, for the most part, the LOESS curves closely followed paths taken by the exponentially smoothed curves served as validation of the method's accuracy in capturing the data's trends and structures, while using ubiquitous smoothing parameters. This comparison also highlighted the disadvantages of exponential smoothing, namely its sensitivity to extreme values and the inability of census tracts to be weighted using the method.

Ultimately, LOESS emerged as the non-parametric smoother of choice, owing to its smoothness and accuracy of data fitness while using ubiquitous smoothing parameters, ensuring speed and consistency in the graphical summarization of metropolitan dynamics and trajectories. The fact that LOESS interpolates explanatory variable values between census tract distances diminishes the detriment of using unstandardized census tract boundaries, and since the graphs illustrate interpolated values between census tracts, creating a continuous gradient emanating from the CBD outwards, it becomes possible to easily analyze changes across time, as opposed to viewing discrete unstandardized boundaries on maps.

Across the social sciences, LOESS is a promising exploratory method for shedding light on the functional forms of non-linear empirical relationships which were previously unknown or uncertain. Metropolitan regions are reflections of the diversities of their inhabitants and how they are spatially distributed across the region. LOESS was demonstrated to be a method providing succinct descriptive snapshots of some of these relationships. Of course, such smoothing methods and the patterns they produce should be viewed as complementary to other methods, including thematic mapping and parametric methods, since parametric and non-parametric smoothers have their own distinctive strengths and weaknesses. However, its flexibility and simplicity makes LOESS a powerful complement to thematic mapping for spatial time-series analyses.

References

Balaban, E., Bayar, A., & Faff, R. W. (2006). Forecasting stock market volatility: Further international evidence*. European Journal of Finance*, 12(2), 171–88.

Beck, N., & Jackman, S. (1998). Beyond Linearity by Default: Generalized Additive Models. *American Journal of Political Science, 42*(2), 596-627. DOI: 10.2307/2991772

Berke, O. (2004). Exploratory disease mapping: kriging the spatial risk function from regional count data. *International Journal of Health Geographics*, 3(1), 18.

Engineering Statistics Handbook - 4.1.4.4. LOESS. (n.d.). Retrieved August 31, 2017, from http://www.itl.nist.gov/div898/handbook/pmd/section1/pmd144.html

Gardner, E. S. (1985). Exponential smoothing: The state of the art. *Journal of Forecasting*, 4(1), 1–28.

Härdle, W., & Marron, J. S. (1995). Fast and simple scatterplot smoothing. *Computational Statistics & Data Analysis*, 20(1), 1–17.

Harrell, F. (2015). *Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis* (2nd ed.). New York, NY: Springer-Verlag.

Hutchinson, M. F. (1995). Interpolating mean rainfall using thin plate smoothing splines. *International Journal of Geographical Information Systems*, 9(4), 385–403.

Martin, K., & Ralphs, M. (2014). Using kernel methods to visualise crime data. *Statistical Journal of the IAOS*, 30(3), 177–83.

Monmonier, Mark. (1996). *How to Lie with Maps.* Chicago: University of Chicago Press.

Jacoby, W. G. (2000). LOESS: A nonparametric, graphical tool for depicting relationships between variables. *Electoral Studies*, 19(4), 577–613.

Peltier, J. (2009, March 09). LOESS smoothing in Excel. Retrieved 2016, from https://peltiertech.com/loess-smoothing-in-excel/

Ravinder, H. V. (2013). Determining the optimal values of exponential smoothing constants – Does solver really work? *American Journal of Business Education* (Oncurve), 6(3), 347.

Rodriguez, G. (2001). *Smoothing and non-parametric regression*. Princeton: Princeton University.

Walks, A. (2011). Economic restructuring and trajectories of socio-spatial polarization in the twenty-first century Canadian city. Chapter in Bourne, L.S., Hutton, T., Shearmur, R. and Simmons, J. (eds.). *Canadian Urban Regions: Trajectories of Growth and Change.* Toronto: Oxford University Press. 125–159.

Walks, A. (2001). The social ecology of the post-Fordist/global city? Economic restructuring and socio-spatial polarisation in the Toronto urban region. *Urban Studies*, 38(3),407–47.