

Establishing an air pollution monitoring network for intra-urban population exposure assessment: A location-allocation approach

Pavlos S. Kanaroglou^{a,*}, Michael Jerrett^b, Jason Morrison^c,
Bernardo Beckerman^b, M. Altaf Arain^b, Nicolas L. Gilbert^d, Jeffrey R. Brook^e

^aCenter for Spatial Analysis, School of Geography and Geology, McMaster University, Hamilton, Ont., Canada L8S 4K1

^bSchool of Geography and Geology and McMaster Institute of Environment and Health, McMaster University Hamilton, Ont., Canada L8S 4K1

^cSchool of Computer Science, Carleton University, Ottawa, Ont., Canada

^dHealth Canada, Air Health Effects Section, Ottawa, Ont., Canada

^eMeteorological Service of Canada, Toronto, Ont., Canada

Received 19 January 2004; received in revised form 21 May 2004; accepted 9 June 2004

Abstract

This study addresses two objectives: (1) to develop a formal method of optimally locating a dense network of air pollution monitoring stations; and (2) to derive an exposure assessment model based on these monitoring data and related land use, population, and biophysical information. Previous studies have located monitors in an ad hoc fashion, favouring the placement of monitors in traffic “hot spots” or in areas deemed subjectively to be of interest. We apply our methodology in locating 100 nitrogen dioxide monitors in Toronto, Canada. Locations identified by the method represent land use, transportation infrastructure and the distribution of at-risk populations. Our exposure assessments derived from the monitoring program produce reasonable estimates at the intra-urban scale. The method for optimally locating monitors may have widespread applicability for the design of pollution monitoring networks, particularly for measuring traffic pollutants with fine-scale spatial variability.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Air pollution; Traffic emissions; Monitoring networks; Location-allocation models; Health effects assessment

1. Introduction

Policymakers and scientists have shown growing interest in the health effects of chronic exposure to ambient air pollution. Recent epidemiological studies have generated specific interest in traffic pollution. For

example, a cohort study from the Netherlands (Hoek et al., 2002) reported that large health effects on cardio-pulmonary mortality were associated with living near a major road (Relative risk = 1.95, 95% CI: 1.09–3.51). Yet uncertainties in exposure assessment methodologies continue to raise questions about the reliability and accuracy of risk estimates from intra-urban air pollution studies (Briggs et al., 2000). In this context, we address two research objectives: (1) to develop a formal method of optimally locating a dense network of air pollution

*Corresponding author. Tel.: +1905 525 9140x24960 or x23525.

E-mail address: pavlos@mcmaster.ca (P.S. Kanaroglou).

monitoring stations; and (2) to derive an exposure assessment model based on these monitoring data and related land use, population, and biophysical information.

This paper represents the initial stage of a multi-year research project funded by the Canadian Institutes of Health Research. The ultimate goal of the project is to assess the relation between traffic-generated air pollution and health outcomes ranging from childhood asthma to mortality from lung cancer. Some of the health data are described elsewhere, and significant associations with background particulate and sulfur dioxide pollutants and elevated mortality have already been reported (Finkelstein et al., 2003). A major limitation of this first health effects study stemmed from relatively crude exposure metrics, which relied exclusively on interpolations from a sparse network of government monitoring stations.

Because field-monitoring studies constitute one of the most expensive components of an epidemiological study, we were interested in developing a method that would make maximal use of scarce air pollution monitoring data. Previous studies have located monitors in an ad hoc fashion, favouring the placement of monitors in traffic “hot spots” or in areas deemed subjectively to be of interest for land use and population characteristics (Lebret et al., 2000; Kukkonen et al., 2001; Goswami et al., 2002). We achieved this first step of our research plan by developing a formal methodology for locating a fixed number of air pollution monitors in Toronto, Canada. Formal methods for the design of pollution monitoring networks have been proposed before (Caselton and Zidek, 1984; Caselton et al., 1992; Silva and Quiroz, 2003; Trujillo-Ventura and Ellis, 1991; Haas, 1992). Our proposed approach is geared for urban applications for air pollution measurements and is designed to capture micro-scale variation, especially around major roadways.

After deriving the monitoring locations, we deployed 100 nitrogen dioxide (NO₂) monitors throughout the City of Toronto. We then implemented a land use regression (LUR) model with the collected monitoring data. For this study, we chose NO₂ as a proxy for traffic pollution because it is relatively inexpensive to measure and has been used widely as a metric of exposure to traffic emissions (Briggs et al., 1997).

2. Methods

2.1. Conceptual framework for locating monitoring stations

The approach we use consists of two stages. First, based on specified criteria, we determine a continuous surface over the study area, which we term “the demand

surface”. Higher values of the demand surface correspond to increased need for monitoring. Second, we use the demand surface as an input to an algorithm that solves a constrained optimization problem from the general family of location-allocation (L-A) problems. The algorithm identifies the optimal locations for a predefined number of air pollution monitors.

We use two criteria for the determination of the demand surface. The first postulates that a larger number of monitors should be located where the pollution surface is expected to exhibit higher spatial variability. To implement this criterion, we need an initial estimate of the pollution surface, which would be an approximation of the expected surface after actual monitoring. Since this initial surface is not generally available, we recommend obtaining a first estimate either with available government monitoring data or with a network of ground stations operated by individual researchers. Both government and individual monitoring networks are usually spatially sparse. In our particular application, we obtained a first estimate of the pollution surface through LUR with data from an area wider than the City of Toronto to encompass a large enough number of fixed-site monitoring stations (i.e., South-Central Ontario). The process is described in more detail in the methodology implementation section below.

Given a first estimate of the pollution surface, spatial variability of pollution $\gamma(\bar{x}, h)$ at location \bar{x} is determined by the following equation, inspired by the semivariogram equation that is often used in geostatistics (Cressie, 1993, p. 58):

$$\hat{\gamma}(\bar{x}, h) = \frac{1}{2} \sum_h \left[\left\{ z(\bar{x}) - z\left(\bar{x} + \bar{h}\right) \right\}^2 \right]. \quad (1)$$

In practice, a grid is imposed on the study area and $\hat{\gamma}$ is calculated at all the centroids of the zones that make up the grid. If $Z(\bar{x})$ represents a random variable of the pollution estimate at location \bar{x} , then $z(\bar{x})$ is a specific pollution estimate at the same location. The right-side summation of Eq. (1) occurs over all possible pairs of locations that are formed between \bar{x} and a location within distance h from \bar{x} . The determination of $\hat{\gamma}$ at all locations \bar{x} provides a surface of variability in the pollution surface. This creates the demand surface that satisfies the first criterion noted above.

To satisfy the second criterion, we appropriately modify the demand surface achieved through Eq. (1). More specifically it might be desirable to intensify pollution monitoring in areas where the density of a population of interest is high. Here, we intensify demand for pollution monitors in areas with high densities of at-risk populations. That population group might have demographic characteristics of importance for the intended health study, such as children or the elderly.

To achieve this effect, a weighting scheme is implemented according to the equation

$$W_R = \frac{P_R/P_T}{\hat{\gamma}_R/\hat{\gamma}_T}. \quad (2)$$

Here P_R is the population of interest in region R within the study area and P_T is the same population for the entire study area. Thus, the numerator represents the proportion of the total population of interest in the study area that resides in region R . Similarly, the denominator represents the proportion of the total variability for the entire study area that can be attributed to region R . The weight W_R is applied to $\hat{\gamma}(\vec{x}, \vec{h})$ for each location \vec{x} that belongs to region R . The weighting scheme is designed so that the ratio of variability estimates for any two locations within region R remains unchanged after weighting. At the same time, the variability for all locations within region R is augmented by the proportion of the population of interest within region R . Furthermore, the total variability for the entire study area is maintained at the unweighted level.

The approaches used in the literature can be classified into two main groups. The first makes use of information theory and devices optimization methods that locate a given number of monitors in a network so that the information content in the collected data is maximized or the uncertainty is minimized (Caselton and Zidek, 1984; Caselton et al., 1992; Silva and Quiroz, 2003). The same idea is applied to evaluating the effectiveness of a government monitoring network of

stations or optimizing the information gain by locating an additional fixed number of new monitors. The second group uses kriging, a geostatistical spatial interpolation technique. The idea is to locate the monitors of a network so that the mean square prediction error (or kriging variance) is minimized (Trujillo-Ventura and Ellis, 1991; Haas, 1992). This method ensures adequate coverage of the study area, since it is well known that the kriging variance is higher in places of sparse sampling.

The basic idea for our approach comes from the statistical stratified sampling theory, whereby strata associated with higher variability are sampled more intensively. The method relies on the identification of the pollution surface variability over the study area. It is designed for capturing the small-scale spatial variability of pollution, especially around expressways in urban areas. As with the other methods, additional objectives in the optimization problem can be added, such as increased sampling in areas of high population density. An advantage of our method is that it relies on well-developed location theory and the associated optimization algorithms, known as L-A.

2.2. The study area and data

For the derivation of an initial estimate surface of pollution over the city of Toronto we used available data from a network of fixed monitoring stations in South-Central Ontario; a densely populated conurbation extending from St. Catharines to Oshawa (roughly 200 km, see Fig. 1) and encompassing Toronto.




Fig. 1. Study area in South-Central Ontario.

Predictors of NO₂ pollution in a LUR model were derived from land use and transportation data acquired from a commercial source, Desktop Mapping Technologies Incorporated (DMTI), Markham, Ont., Canada. Population data at the census tract (CT) level were obtained from the standard tabulations of census data provided by Statistics Canada.

Toronto, the study area, is Canada's largest city with an estimated population of 4.7 million people (Statistics Canada, 2001) and an approximate area of 633 km². Located on the north shore of Lake Ontario (43° 39'N, 79° 23'W), it is classified as being in the "Humid East" region of temperate North America (Getis and Getis, 1995). Similar to other large cities in North America, many expressways traverse the Toronto landscape, including some of the busiest in North America. For example, Highway 401 has an estimated flow of about 400,000 vehicles per day.

Monitor deployment in Toronto, based on the locations derived with the method we describe in this paper, occurred in the time period 9–25 September 2002, at 100 locations. OgawaTM passive samplers were used to measure concentrations of NO₂. We deployed two-sided samplers in pairs of two (yielding four observations per location) at a height of 2.5 m. All samplers were removed 14 days after their installation. We used ion chromatography to determine the nitrite content on

collection pads (Gilbert et al., 2003). This field sampling yielded 95 valid measurements. Five samples were lost due to vandalism or equipment malfunction.

2.3. Estimating the initial pollution surface as input for the location-allocation model

To estimate the initial pollution surface for determining the spatial variability, we employed a LUR. LUR uses pollution concentrations as the dependent variable and proximate land use, traffic, and physical environmental variables as independent predictors. This methodology thus seeks to predict pollution concentrations at a given site based on surrounding land use and traffic characteristics. Specifically, this method uses measured pollution concentrations (y) at locations (s) as the response variable and land use types (x) within circular areas around s (called buffers) as predictors of the measured concentrations (see Fig. 2). The incorporation of land use variables into the interpolation algorithm detects small-area localized variations in air pollution more effectively than standard methods of interpolation such as kriging (Briggs et al., 1997, 2000; Lebret et al., 2000).

Mean NO₂ estimates at 16 locations shown in Fig. 1, operated by the Ontario Ministry of the Environment (MOE), served as the dependent variable in the LUR.

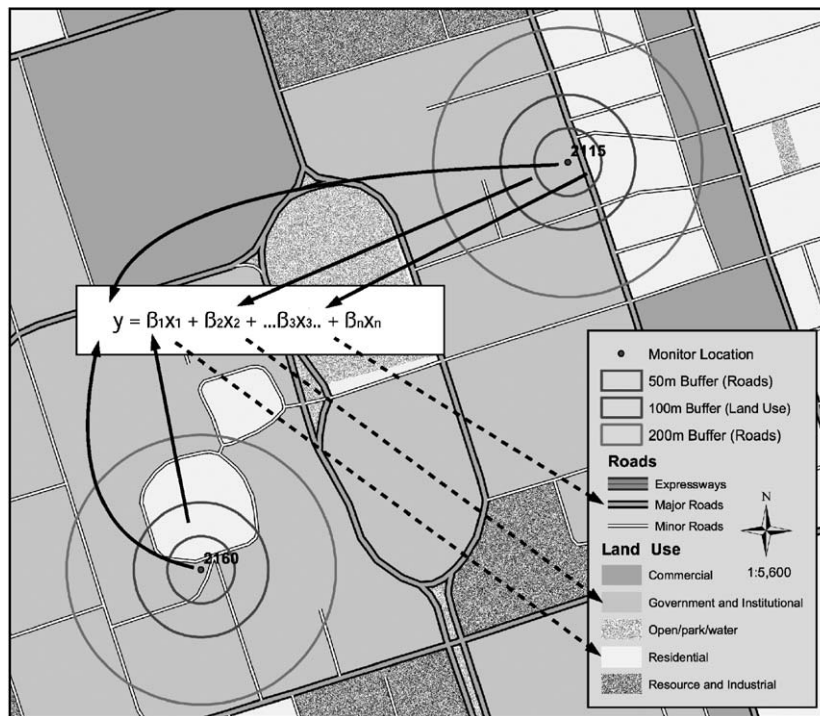


Fig. 2. Elements of a land use regression model showing monitoring locations for NO₂ as the response variable and land use characteristics within buffers as the predictor or independent variable.

The 1999 NO₂ annual mean values at the 16 MOE air monitoring stations in the study area range between 12.4 and 28.4 parts per billion (ppb). Independent variables were obtained by measuring the respective area and length parameters within circular neighbourhood buffers centred on the MOE station position. Accurate georeferencing was essential for the locations of pollution monitoring sites, roads and houses due to the highly localized nature of the estimates. We used a global position system to mark all the monitoring sites (ground-level accuracy of 7 m or less). Field validation of the DMTI land use coverage revealed high accuracy in attribute classification and spatial coordinates.

Land use variables were calculated in hectares and road lengths in kilometres. These variables were measured under a circular buffer that extends from each monitoring location out to a radius of 100 m. Road length parameters were measured under two separate buffers. A first circular buffer extends from each monitoring point out to a radius of 50 m. A second buffer is in the shape of an annulus (donut), with the inner edge at a radius of 50 m and the outer at 200 m. All area and length calculations were performed using ArcView 3.2 software's Spatial Analyst extension (ESRI Corp., Redlands, CA, USA).

In total, 16 land use and transportation variables were tested in a manual stepwise regression with S-Plus 6 software (Mathsoft, Cambridge, MA) to obtain a best estimate pollution model. A bivariate ordinary least-squares regression analysis was used in the initial step to develop a model that predicts the best pollution surface. The variable selection criterion was set at $p < 0.25$ statistical significance for subsequent consideration of the variable (Younger, 1985). The most statistically significant variables were sequentially included until a parsimonious model was achieved. While most variables included in the model achieved statistical significance at conventional levels (i.e., $p < 0.05$), we relaxed this criterion if inclusion of the variable eliminated other problems with outliers.

The resulting estimated LUR model was subsequently used to predict NO₂ values for all cells of a grid imposed over Toronto at 5 m resolution. For a given grid cell, values for independent variables of the LUR equation were calculated within the same neighbourhood around the centre of the grid cell as they were initially calculated for inclusion in the regression equation. If in the regression equation, for example, a significant independent variable is the length of expressway portions that fall within 50 m of a NO₂ measuring station, then for the prediction of NO₂ for a given grid cell we construct a 50 m radius circle around the centre of the grid cell and calculate the length of the expressways within that circle. The 5 m resolution was chosen to allow a reasonable approximation of the vector dataset into raster format. The result of this exercise was a first approximation of a

pollution surface that is subsequently used to derive a monitoring demand surface. Our first task towards this end is to calculate the spatial variability of the pollution surface at every single grid cell, as described in the next subsection.

2.4. Creating a surface of monitoring demand

Air pollution over a study area is spatially continuous. Consistent with geostatistical modelling, we assume that any pollution measurement $z(\bar{x})$ at location \bar{x} is a specific instance of a random variable $Z(\bar{x})$ at the same location. The random variable Z over locations in the study area represents the spatial stochastic process under study. Because of the presence of second-order effects (autocorrelation), we expect the correlation of any two such random variables to diminish with distance, while the variance will increase (Webster and Oliver, 1990). Previous results of NO₂ monitoring (Hewitt, 1991; Gilbert et al., 2003) suggest that little correlation exists if the distance between any two points $|h|$ is "sufficiently large". Also, while correlation varies according to topography and meteorology, 80–90% of traffic related effects come from less than 150–300 m away (English et al., 1999; Briggs et al., 2000). Thus, we select as the maximum distance within which to examine local variability of pollution to be 300 m.

On the basis of these observations, a suitable operational estimator of local variability of NO₂ pollution at location \bar{x} is obtained from Eq. (1) by fixing the maximum distance $|h|$ to be equal to 300 m:

$$\hat{\gamma}(\bar{x}) = \frac{1}{2} \sum_{|\bar{h}| \leq 300 \text{ m}} \left(z(\bar{x}) - z(\bar{x} + \bar{h}) \right)^2. \quad (3)$$

It is worth noting that in Eq. (3), location \bar{x} is represented by the centre of a 5 m grid cell. Variability for a given grid cell \bar{x} is calculated by applying the summation in the right-hand side of Eq. (3) over the pairs that are formed between \bar{x} and all the cells within 300 m from \bar{x} . By contrast, for a semivariogram estimator, the summation in the right-hand side is over all location pairs that are a certain lag $|h|$ apart, and the calculation is performed over several lags to determine the semivariogram function. Thus, while Eq. (3) resembles the standard semivariogram equation, its usage in the present context is different. To determine the surface of the monitoring demand, Eq. (3) is applied for all locations \bar{x} within the study region R .

2.5. Augmenting demand for socio-demographics

The demand intended for sampling stations computed via Eq. (3) accounts for the need to represent pollution variability. When deriving pollution exposures for epidemiological studies, however, it may be desirable

to measure more intensively in areas where the population possesses particular socio-demographic (SD) characteristics of interest. To achieve area-specific weighting of the sampling network, we aim at increasing demand in areas with high concentration of the target populations and decreasing it in areas of low concentration.

We start with the density of the population of interest at the CT level making use of standard tabulation data provided by Statistics Canada. We chose the CT level because it provides a compromise between spatial detail and area homogeneity. Statistics Canada's CTs are equal in size to neighbourhood-like areas, having approximately 2500–8000 people (average population ~4000). CT boundaries generally follow permanent physical features, such as major streets and railway tracks, and attempt to approximate cohesive socio-economic areas, but may not do so because of subjectivity in unit selection or neighbourhood change over time.

We transform the CTs population densities into population counts at the level of grid cells with 5 m resolution. We then aggregate the population counts to grid cells of 2500 m resolution. This is done to account for the spatial subjectivity built into the CT unit. A further advantage of re-aggregating the CT-derived population surface is that the weighted surface is unlikely to have gaps in areas where the population is low. The re-aggregation thus ensures that the final demand surface represents both the population density and pollution variability.

The weighting was implemented as a bivariate linear rescaling of the pollution variability surface. This process conserved the total amount of demand in the initial variability surface. The exact formula we used is an adaptation of Eq. (2) as follows:

$$W_{2500\text{ m}} = \frac{P_{2500\text{ m}}/P_T}{\hat{\gamma}_{2500\text{ m}}/\hat{\gamma}_T}, \quad (4)$$

where $P_{2500\text{ m}}$ is to population of a grid cell at 2500 m resolution and P_T is the total population of the study area calculated as the sum of populations for all grid cells at the 2500 m resolution. $\hat{\gamma}_{2500\text{ m}}$ is the variability in a 2500 m resolution grid cell. For the purposes of this paper, the population at risk was children 6 years of age or younger. By multiplying the derived $W_{2500\text{ m}}$ value with the variability value for all grid cells, a modified demand surface is derived. In subsequent health studies, this exposure assessment will assist with testing the association between onset of childhood asthma and traffic pollution in a case-control design.

2.6. Computing monitoring locations using demand

The discussion of computing demand does not address the issues of calibration, reliability or validation of the NO_2 measurements. For purposes of reliability, it

is advisable to deploy two or more passive samplers at a single location; the measured value at a location is taken to be an average of the measurements at that location. A further level of accuracy is added to our network by calibrating the measurements against the collected MOE data using continuously operating active monitors. This, however, entails that the MOE stations be used as sampling locations in our network.

The remainder of this section details a technique for computing the locations of n number of monitoring stations, assuming that n is a specified constant before the computation. The computation of the locations begins by calculating the locational demand for a grid at 500 m resolution. For Toronto, the centroid of each of the grid cells results in a lattice of 2537 potential locations, which we refer to as "candidate locations".

The goal of our L-A procedure is to select n primary station locations taking into account the demand surface for monitoring that we have created. In formulating this problem as a classic facility location problem, we constrain ourselves to problems solved using the ESRI's ARC/INFO software. This software offers two options: the P-Median Problem or P-MP (ReVelle and Swain, 1970) and the Attendance Maximizing Problem or AMP (Holmes et al., 1972). For our particular problem, we have 2537 candidate locations for 100 monitors to be located ($n = 100$). The candidate locations also act as the demand locations. Each demand location is weighted by the value of the demand surface in that location.

The P-MP places the 100 primary stations in a way that the sum of weighted distances for all demand locations from their nearest station is minimized. The AMP, on the other hand, locates stations so as to maximize an objective function of the form

$$Z = \sum_{i=1}^k \sum_{j=1}^m w_i (1 - bd_{ij}) x_{ij}, \quad (5)$$

where k is the number of demand locations and m is the number of candidate locations. In our case $k = m = 2537$. The weight w_i at location i represents demand, while d_{ij} is the distance between locations i and j . x_{ij} is the allocation decision variable attaining the value of 1 if demand location i is served by a station in j and 0 otherwise. Attendance linearly decreases with distance at the rate of parameter b , the value of which is determined by the particular problem in hand. For our case, we assume that attendance becomes negligible or 0 for a distance of 1.5 km, thus b is determined by solving the Eq. (1) $-b \cdot 1.5 = 0$, which leads to $b = 0.67$.

The above discussion suggests that the AMP formulation provides better control and is thus better suited than the P-MP formulation for our particular problem. The assumption that attendance becomes zero for a distance beyond 1.5 km is justified on the basis of previous research. Findings suggest that concentrations

of NO₂ and other pollutants decrease away from polluting sources (English et al., 1999; Gilbert et al., 2003). The measured decay of concentration around Canadian highways is such that no correlation exists between the highway and a point 300 m away from the highway. For the purposes of the model we assumed this distance to be 1500 m. The definition of the allocation decision variable x_{ij} above provides one of the five constraints of the constrained optimization AMP. Other constraints restrict the number of stations to be located to exactly $n = 100$, and ensure that no two stations are located in the same place.

The actual computation was performed using the L-A functionality of ESRI's ArcPlot module of ArcGIS. We set the optimization criteria to be the "Maximum Attendance (MaxAttend)" problem, and then ran the L-A function. Near-optimal or "heuristically optimal" solutions to these types of computationally intensive problems can be found in a reasonable amount of time using proven heuristic techniques. The Global Regional Interchange Algorithm (GRIA) (Densham and Rushton, 1992) is appropriate for large datasets because running time is linear with respect to n , which allows for computation with current desktop computing technology. The grid resolution affects the optimality of the solution. Given $n = 100$, the choice of a 500×500 m² grid of demand locations made for a problem of reasonable size for desktop computers that could be easily replicated by others.

3. Results

3.1. Monitoring location selection

Results of the initial LUR are displayed in Table 1 and include regular diagnostic values such as sum of

squares (SS), degrees of freedom (df), mean square (MS), and a inter-variable collinearity factor called the various inflation factor (VIF). The table shows that commercial, government and institutional, expressway (0–50 m), and expressway (50–200 m) are significant land-use and transportation variables in predicting NO₂ with the available dataset. Similar to other LUR results (Briggs et al., 2000), some coefficients take an unexpected sign (i.e., 0–50 is negative), but this appears to balance the other positive signs in the equation, and this model produced maximal prediction with minimal bias (measured by the Mallows' Cp statistic) compared to other models we tested. The pollution surface generated from this model had high pollution concentrations predicted in the vicinity of the expressways and relatively low pollution predictions in other areas.

Modelled pollution values near expressways ranged from 29.5 to 1502 ppb, while in other areas the pollution estimates ranged from 18.1 to 34.5 ppb. Given the range of the MOE NO₂ monitoring concentrations and those noted in other studies (Hewitt, 1991), such high levels are unrealistic. Yet, independent of the actual values, an important characteristic was preserved in the surface, namely the simple relative pollution variability caused by the different land uses—an indication that some areas have higher pollution variability than others. The variability was obtained by Eq. (3) using the obtained pollution surface. The obtained surface of variability appears to have much the same characteristics as the pollution surface, with high values near expressways and low values in other areas.

Proximity to expressway features dominated sampling locations selected by the L-A model, as seen in Fig. 3. About 95% of the stations were assigned to areas that coincided with expressway features in the pollution model. The other 5% of the locations satisfied the

Table 1
Initial land use regression model used for location allocation procedure to locate samplers

Source	SS	df	MS		
Regression	213.89	4	53.47	Number of obs = 16	
Residual	130.68	11	11.88	$F(4,11) = 4.5$	
Total	344.58	15		Prob > F = 0.021	
				$R^2 = 0.621$	
				Adj. $R^2 = 0.483$	
				Prob > t	
Variable	Coefficient	Std. error	t	VIF	
NO ₂ (ppb)					
Constant	18.058	1.353	13.348	0.000	
Commercial	5.223	2.080	2.511	0.029	1.0
Government/institutional	1.748	1.065	1.641	0.129	1.1
Expressway (0–50 m)	–592.725	228.883	–2.590	0.025	2.3
Expressway (50–200 m)	253.155	68.532	3.694	0.004	2.4

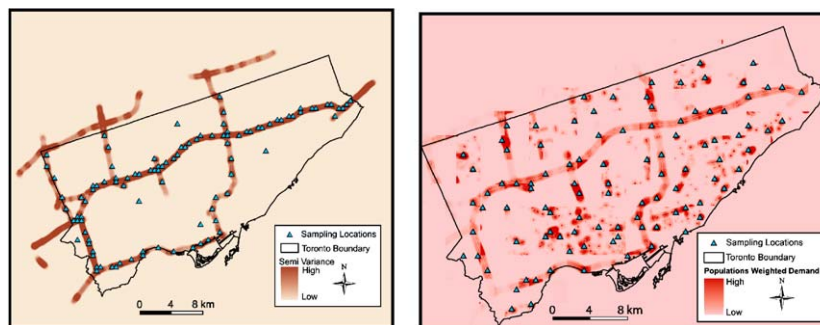


Fig. 3. Comparison of unmodified sampling locations to population weighted locations.

attendance demand created by the surrounding expressways and consequently were located centrally.

In contrast, the population weighting method (PWM) produced a well-distributed monitoring network, as shown in Fig. 3. There is a low-level clustering in the central west part of the city, an area with the highest aggregation of child populations.

Once the L-A procedure was complete and alternative monitoring networks were configured, it was necessary to ascertain the monitoring networks' capacity to represent key land use and transportation criteria. Table 2 presents a comparison of the network developed by taking into account only the pollution surface variability criterion with the network that also took account of the population distribution criterion (children of age 0–6). The comparison of the networks is based on ten land-use criteria that were used in the LUR model, as shown in the first column of the Table. The second and third columns are based on characteristics of all the 2537 candidate locations. They indicate that 245 of the candidate locations are between 50 and 200 m away from an expressway, while 96 of them are within 50 m from an expressway. Similarly, the pollution variability criterion alone gives rise to a network whereby 94 of the 100 stations are located between 50 and 200 m away from an expressway as compared to 32 of the 100 stations located within the same distance band from an expressway in the case where the population criterion is also taken into account. This indicates that the variability criterion alone produces a much more concentrated network pattern around the expressways.

This conclusion is strengthened by the observation that in the variability criterion case 58 of the 100 stations are located within 50 m of an expressway, compared with 32 stations in the case where population is taken into account. Similarly, the population-weighted surface allocates more stations in residential areas and fewer stations in industrial areas. The latter is probably because expressways go through industrial areas. Comparing the mean distance values of the two

networks with those of all the candidates reveals that in six out of ten land use cases, the means are significantly different for the variability criterion, but in only two land use cases the difference in means is statistically significant when population weighting is used. We thus conclude that the population-weighted surface produces a better representation of the land use and transportation characteristics of the study area.

3.2. Results from the Toronto deployment

Of the 100 monitors deployed in the Toronto area, 95 produced valid measurements. These locations were used in a second LUR model to derive a predicted surface of NO₂. The arithmetic mean of NO₂ values (measured in ppb) for these 95 records was 32.71 ppb, with values ranging from 17.41 to 77.31 ppb, which is larger than the value we observed through the government-monitoring network. In total, 79 land use, transportation, population, and physical geography variables were tested (see Jerrett et al. 2003 for more detailed descriptions of the variables).

Logarithmic NO₂ concentrations were negatively correlated with the distance with the nearest expressway and the area of open space within 400 m, while they were intuitively correlated with road measures and dwellings in the vicinity. The final regression model (shown below) yielded a R^2 of 0.633 (see Table 3).

These results confirm that in addition to contributing to regional air pollution, urban traffic is a strong determinant of intra-urban variation in air pollution. Results also indicate that a combination of traffic and land use variables could be used to estimate exposure to traffic-related air pollution in epidemiologic studies on the effects of air pollution on health. Unlike the initial model based on the sparse monitoring network, which we used to derive the variability of the initial pollution surface, all the coefficients take the expected sign. This LUR model also resulted in a predictive pollution surface with the expected characteristics (see Fig. 4) with

Table 2
Land use and transportation criteria represented by location-allocation models compared to candidate location values

<i>Mean std. deviation</i>	Candidate	<i>N</i>	Variability criterion (VC)	<i>N</i>	VC+ popul age 0–6	<i>N</i>
Distance	1783.6 1381.2	2537	158.5 496.1	100	1631.7 1620.2	100
Expressway 50–200	263.2 190.4	245	417.3 189.3	94	380.4 207.7	32
Express 0–50	44.8 29.7	96	52.2 32.5	58	50.2 30.7	20
Major 50–200	90.6 45.3	1151	91.9 46.0	58	100.5 49.9	63
Major 0–50	19.0 8.6	344	18.3 9.9	21	17.4 9.2	21
Minor 50–200	229.3 119.6	2264	128.8 93.5	76	200.4 107.0	93
Minor 0–50	24.1 10.7	1474	18.3 10.6	30	22.9 10.1	52
Residential 100	856.8 425.3	1779	541.5 439.2	53	672.8 410.1	73
Industrial 100	660.8 459.2	774	594.0 403.5	45	649.3 453.8	36
Commercial 100	375.8 349.3	235	353.6 311.8	10	481.1 402.0	39
Gov. Inst. 100	415.7 394.0	434	659.3 469.7	10	451.1 435.1	19

only slight over prediction near the expressways. Areas in proximity to expressways and in the downtown core appeared to have higher levels of NO₂, while areas with less development in the northeast of the city exhibit lower levels.

4. Conclusion

With the increasing need to understand the spatial distribution of air pollution at the intra-urban scale, the methods developed here for locating air pollution monitors have proven to be an effective and systematic technique to maximize sampling coverage in relation to important socio-demographic (SD) characteristics and likely pollution variability. The location-allocation (L-A) approach offers flexibility to integrate an assortment of variables into the demand surface that can reconfigure the resulting monitoring network. In this application, the variable integration has been implemented with a SD weighting of children 6 years of age and under. Because we had to rely on sparse initial pollution data, the population-weighted pollution surface also ensured

that the subsequent sampling locations would yield useful results for human exposure assessment.

To validate this method, we plan to replicate the entire modelling process using the pollution surface generated from the 95 NO₂ observations collected in Toronto from our first round of monitoring. Assessing the impact of availability of better pollution data on the monitoring locations might have important implications for the generalizability of this method. If we achieve significantly different results from an improved input pollution surface, then this L-A approach would be effectively limited to the few places where spatially extensive air pollution data are available or to the second round of monitoring. If, however, the model performs adequately with weaker pollution data combined with SD population data, then the method may have wide applicability to most urban areas. Given the large influence of population weighting in our findings and the subsequently reasonable results from the first round of monitoring in Toronto, the method appears to have considerable promise for improving the assessment of exposure to ambient air pollution in epidemiologic studies.

Table 3
Land use regression model created with 95 sampled data values

Source	SS	df	MS		
Regression	5.195	6	.866	Number of obs = 95	
Residual	3.009	88	.034	$F(6,88) = 25.318$	
Total	8.204	94		Prob > $F = 0.000$	
				$R^2 = 0.633$	
				Adj. $R^2 = 0.608$	
				Prob > t	
Variable ^a	Coefficient	Std. error	t	Prob > t	VIF
LN(NO ₂)					
(Constant)	3.460	.087	39.709	.000	
DIST_EX	-5.291E-05	.000	-3.554	.001	1.564
RD1_200	.197	.030	6.511	.000	2.058
RD3_3005	-4.078E-02	.015	-2.733	.008	2.101
RD2_500	6.949E-02	.019	3.664	.000	1.350
OPEN400	-1.036E-02	.003	-3.306	.001	2.265
DC2000	6.950E-05	.000	3.161	.002	1.652

^aRD1_200—measure of expressway within 200 m, RD3_3005—measure of local roads within a donut-shaped area with an inner radius of 300 m and outer radius of 500 m, RD2_500—measure of major roads within 500 m, OPEN400—measure of park, open, recreational, or water body land within 400 m, DC2000—density of dwellings within 2000 m (Kernel estimate), DIST_EX—distance to the nearest expressway.



Fig. 4. Land use regression model of NO₂ concentrations (ppb) using 95 sampled values.

Acknowledgments

We are grateful to Health Canada and the Canadian Institutes of Health Research for funding this work. We also thank Pat De Luca, Dan Crouse, Norm Finkelstein, and Chris Giovis for assistance with the GIS modelling. The first author is grateful to the Social Sciences and Humanities (SSHRC) Canada Research Chairs (CRC) program.

References

- Briggs, D., Collins, S., Elliott, P., Fischer, P., Kingham, S., Lebre, E., et al., 1997. Mapping urban air pollution GIS: a regression-based approach. *International Journal of Geographical Information Science* 11, 699–718.
- Briggs, D., de Hoogh, C., Gulliver, J., Wills, J., Elliott, P., Kingham, S., et al., 2000. A regression-based method for

- mapping traffic-related air pollution: application and testing in four contrasting urban environments. *Science of the Total Environment* 253, 151–167.
- Caselton, W.F., Zidek, J.V., 1984. Optimal network monitoring designs. *Statistics and Probability Letters* 2, 223–227.
- Caselton, W.F., Kan, L., Zidek, J.V., 1992. Quality data networks that minimize entropy. In: Guttorp, P., Walden, A. (Eds.), *Statistics in Environmental and Earth Sciences*. Griffin, London.
- Cressie, N., 1993. *Statistics for Spatial Data*. Revised edition. Wiley, New York.
- Densham, P., Rushton, G., 1992. A more efficient heuristic for solving large p-median problems. *Papers in Regional Science* 71, 307–329.
- English, P., Neutra, R., Scaif, R., Sullivan, M., Waller, L., Zhu, L., 1999. Examining associations between childhood asthma and traffic flow using a geographic information system. *Environmental Health Perspectives* 107, 761–767.
- Finkelstein, M., Jerrett, M., DeLuca, P., Finkelstein, N., Verma, D., Chapman, K., Sears, M.R., 2003. Relation between income, air pollution and mortality: a cohort study. *Canadian Medical Association Journal* 169, 397–402.
- Getis, A., Getis, J., 1995. *The United States and Canada: The Land and the People*. Wm. C. Brown Publishers, Dubuque, IA, USA.
- Gilbert, N.L., Woodhouse, S., Stieb, D.M., Brook, J.R., 2003. Ambient nitrogen dioxide and distance from a major highway. *The Science of the Total Environment* 312, 43–46.
- Goswami, E., Larson, T., Lumley, T., Liu, L.-J.S., 2002. Spatial characteristics of fine particulate matter: identifying representative monitoring locations in Seattle. *Washington. Journal of Air & Waste Management Association* 52, 324–333.
- Haas, T.C., 1992. Redesigning continental-scale monitoring networks. *Atmospheric Environment* 26A (18), 3323–3333.
- Hewitt, C.N., 1991. Spatial variations in nitrogen dioxide concentrations in an urban area. *Atmospheric Environment* 25B, 429–434.
- Hoek, G., Brunekreef, B., Goldbohm, S., Fischer, P., van den Brandt, P.A., 2002. Association between mortality and indicators of traffic-related air pollution in the Netherlands: a cohort study. *Lancet* 360 (9341), 1203–1209.
- Holmes, J., Williams, F.B., Brown, L.A., 1972. Facility location under maximum travel restriction: an example using daycare facilities. *Geographical Analysis* 4, 258–266.
- Jerrett M., Arain, M.A., Kanaroglou, P., Beckerman, B., Crouse, D., Gilbert, N.L., Brook, J.R., Finkelstein, N., 2003. Modelling the intra-urban variability of ambient traffic pollution in Toronto, Canada. *Proceedings: Air-Net Conference, Rome Italy, 2003*.
- Kukkonen, J., Harkonen, J., Karppinen, A., Pohjola, M., Pietarila, H., Koskentalo, T., 2001. A semi-empirical model for urban PM10 concentrations, and its evaluation against data from an urban measurement network. *Atmospheric Environment* 35, 4433–4442.
- Lebret, E., Briggs, D., van Reeuwijk, H., Fischer, P., Smallbone, K., Harsemma, H., Kriz, B., Gorynski, P., Elliott, P., 2000. Small area variations in ambient NO₂ concentrations in four European areas. *Atmospheric Environment* 34, 177–185.
- ReVelle, C.S., Swain, R.W., 1970. Central facilities location. *Geographical Analysis* 2, 30–42.
- Silva, C., Quiroz, A., 2003. Optimization of the atmospheric pollution monitoring network at Santiago de Chile. *Atmospheric Environment* 37, 2337–2345.
- Statistics Canada, 2001. 2001 Census of Canada. Available online: <http://www12.statcan.ca/english/census01/release/index.cfm>.
- Trujillo-Ventura, A., Ellis, J.H., 1991. Multiobjective air pollution monitoring network design. *Atmospheric Environment* 25A (2), 469–479.
- Webster, R., Oliver, M., 1990. *Statistical Methods in Soil and Land Resource Survey*. Oxford University Press, Oxford.
- Younger, M.S., 1985. *A First Course in Linear Regression*, Second ed. PWS Publishers, Boston, Mass.