ENHANCING THE DESIGN OF OBSERVATIONAL STUDIES OF COMMUNITY POLICING: USING GEOSPATIAL DATA MINING TO DESIGN NON-EXPERIMENTAL PROGRAM EVALUATIONS

Christian KREIS

Abstract

The current research is an application of geospatial data mining algorithms to enhance the validity of an observational study of community policing in Switzerland's major urban areas. Both unsupervised and supervised data mining algorithms are used to cluster high-dimensional data on neighbourhood-level crime rates, the socio-economic and demographic structure, and the built environment in order to identify matched comparison areas across the five cities for the subsequent impact evaluation. The resulting neighbourhood typology reduced the within-cluster variance of the contextual variables and accounted for a significant share of the between-cluster variance in the survey measures of community policing impact. This suggests that geo-computational methods help to balance the observed covariates and hence to reduce threats to the internal validity of a non-experimental research design. The assessment of the validity of the neighbourhood classification system for evaluation purposes and its geo-visualization for better communication with practitioners and intelligence-based decision making form an integral part of the study.

Keywords

neighbourhood profiling, impact evaluation, machine learning, geospatial data mining

Résumé

La présente étude est une application d'algorithmes d'exploration de données géo-spatiales afin d'améliorer la validité d'une étude observationnelle de la police de proximité dans les plus grands centres urbains de Suisse. Des algorithmes supervisés et non supervisés ont été utilisés sur les données à haute dimensionnalité, relatives à la criminalité à l'échelle des quartiers, à la structure socio-économique et démographique et au cadre bâti. Le but poursuivi est d'identifier des zones comparables à travers les cinq villes étudiées afin d'analyser les impacts de la police de proximité. La typologie de quartier développée a abouti à une réduction de la variance intra-groupe des variables contextuelles. Elle permet d'expliquer une partie significative de la variance intergroupe des indicateurs d'impacts de la police de proximité recueillis par le biais de sondages. Ceci semble suggérer que les méthodes de géo-informatique aident à équilibrer les co-variables observées et donc à réduire les menaces relatives à la validité interne d'un concept de recherche non-expérimental. L'analyse de la validité de la typologie des quartiers à des fins d'évaluation ainsi que sa géo-visualisation pour une meilleure communication des résultats aux praticiens et l'aide à la décision stratégique font partie intégrante de l'étude.

Mots-clés:

criminologie environnementale, prise de décision, inférence, auteurs en série, profilage géographique, analyse spatio-temporelle

I. INTRODUCTION:

A recurrent demand in recent years in the area of crime prevention has been that programs be "evidence-based", meaning that criminal justice policies should be subjected to scientific evaluation in order to identify best practices (Sherman et al., 2002). The methodological standards of scientific program evaluation in general have been cogently defined already during the 1960s and 1970s, and have been reaffirmed more recently with particular reference to criminological interventions (Cook & Campbell, 1979; Shadish et al., 2002; Farrington, 2003). This body of knowledge posits a clear hierarchy of the methodological quality of different research designs, with the randomized controlled trial (RCT) held as the "gold standard" of

scientific evaluation.

In the areas of crime prevention and policing, however, RCT designs have seldom been implemented because field experiments are deemed as politically risky, or ethically questionable, or both. For area-based criminological interventions targeted at specific places or entire jurisdiction, finding a sufficient number of treatment and control areas can be challenging and statistical power correspondingly low. As a result, several authors have bemoaned a dearth of methodologically sound program evaluations (e.g. Weisburd & Eck, 2004; Welsh & Hoshi, 2002).

An observational study is the alternative empirical analysis of the effects of a treatment intervention in cases where an experimental design is either unethical or infeasible. A good observational study strives to emulate the key aspects of a RCT design in order to enhance the validity of its conclusions. Crucially, in a true experiment, the distribution of covariates is similar between treatment and control group as a result of the random assignment of the study objects to the treatment and control condition. An observational study seeks to achieve this by selecting a set of comparisons that resemble the treated objects on the observable covariates prior to the treatment intervention. Matching techniques are then used to achieve a similar distribution of the observed covariates (though not the unobserved covariates) between the treated objects and the selected controls (Rosenbaum, 2010).

The current research forms part of an observational study of community policing in major Swiss urban areas (Kreis, 2012). Community policing is both a philosophy and an organizational strategy of the police that promotes a renewed partnership between the police agency and local communities to solve problems of crime and disorder. For the current study, the selection of suitable treatment and control areas for the planned evaluation was compounded by the fact that police forces in Swiss cities, beginning in the late 1990s, rapidly introduced community policing across their entire jurisdiction without making any provisions for later evaluation. In the current context this meant that any valid control group had to be found outside each urban area and baseline data for any pre-test/post-test comparisons had to come from existing data sources.

The exploratory data analysis, which had been undertaken as a preliminary study (Kreis, 2009), revealed that the spatiotemporal patterns of the four theoretical constructs of community policing impact – crime, fear of crime, neighbourhood disorder and public attitudes towards the police – displayed some remarkable parallels across the five urban areas. In particular, the exploratory spatial analyses established that the patterns of crime rates and perceptions of disorder had remained rather stable over the short and medium run, whereas areas with elevated levels of fear had shifted from the urban centres to the city boundaries between the late 1980s and 2005. Moreover, whereas these observable

response patterns were noticeably different between the Swiss German and Swiss French cities, responses within a given language region proved to be unexpectedly homogenous (Kreis, 2012).

These rather systematic spatiotemporal patterns of the outcome indicators implied that an impact evaluation of community policing over an extended study period that did not control for shifting neighbourhood characteristics, risks being unreliable at best and positively misleading at worst. The striking parallels between the cities under study gave rise to the idea of developing a neighbourhood typology to match similar neighbourhoods across urban areas in order to study the impact of different community policing strategies in similar neighbourhood contexts. The objective is thus to develop a classification system in order to group neighbourhood areas into clusters of similar type based on a series of environmental, socio-economic and socio-demographic indicators. This approach is based on the premise that the spatial dynamic of the socio-economic processes unfolding in a city affect the crime and response patterns to a considerable extent and that these processes would repeat themselves from one city to another.

II. THEORETICAL CONSIDERATIONS

A. The rationale of matching neighbourhood areas for performance evaluation

The classification or profiling of neighbourhoods or bigger administrative areas in the field of law enforcement and policing has been tried and applied primarily in England and Wales in an effort to increase police accountability and to set benchmarks to measure and improve police performance. In the mid-1990s, the government police inspectorate thus created the *most similar force* group that assigned all 43 separate police forces into groupings of similar type (Ashby & Longley, 2005, 56f.). In the early 2000s, the British Home Office published similar groupings of smaller scale police administrative units, which classified the more than 300 Crime and Disorder Reduction Partnerships (CDRP) and Basic Command Units (BCU) across England and Wales into *families* of similar type (Sheldon et al., 2002). The rationale behind the clustering of policing units was to identify areas that faced similar *policing environments* and were thus suited for meaningful cross-sectional comparisons to evaluate performance (Ashby & Longley, 2005, 56f.).

The classification and matching of police administrative units for performance evaluation is based on the premise that different areas differ significantly in their responsiveness to different policing styles. It rests on the observation that even though crime and poverty are correlated, not all deprived areas are equally crime-ridden. The clustering to categorize the different poli-

cing areas therefore must include not only ecological characteristics of an area such as the socio-economic status (SES) or demographic composition but data on attitudes and lifestyles as well. The very importance of such *soft* attitudinal aspects has been underlined by analyses of the British Crime Survey (BCS) data from the 1990s, which showed that even though actual levels of victimization had been falling, two thirds of respondents were under the impression that crime had gone up. This apparent mismatch between falling crime levels and the widely held belief of an increase in crime has become known in the literature as the *reassurance gap* and has spawned *reassurance policing*, which aims both to rectify the public's perception and ultimately to provide safer neighbourhoods (Williamson et al., 2006, 191-4).

Linking British census and BCS data, Williamson et al. (2006) developed an indicator of an area's level of social capital and compared these values to actual victimization rates. Their results showed not only that areas with higher levels of social capital suffered comparatively lower levels of victimization but also that people's perception of crime, crime reporting, fear of crime and attitudes towards the police differed according to the composition of their neighbourhood. The authors thus concluded that such geo-demographic profiling serves as a useful tool to design reassurance policing or community policing strategies, which are more likely to be effective if targeted specifically at the needs of each type of neighbourhood area.

Ashby and Longley (2005, 427-32) proposed three kinds of geo-demographic analyses that may support police strategic decision-making and performance evaluation: area profiling, operational data profiling and crime survey profiling. Firstly, the basic profiling of police patrol beats or precincts into neighbourhoods of different types provides basic strategic intelligence. Mapping such a typology in a GIS provides an additional spatial dimension to this kind of information. Secondly, the profiling of crime events and police operational data allow police analysts to compute the propensities of specific crime events in different neighbourhood types. Such information makes it possible to identify areas with unexpectedly high or low levels of victimization or to assess the effects of targeted policing interventions. Finally, adding survey data to the analysis helps unearth likely variations in popular attitudes to disorder, fear of crime and the police across different neighbourhood types. If the place of residence of each respondent is known, survey data can be pooled by neighbourhood type to calculate area-level, regional, or even national scores, which may then be extrapolated for analysis at the local level.

The current study builds on and tests these theoretical arguments for a comparative evaluation of community policing across Switzerland's five biggest urban areas. In order to match similar neighbourhoods across the five cities, the current study aims to create a classification system of urban neighbourhoods, which in few dimensions aptly describes the spatiotemporal patterns observed in the high-dimensional input data. The idea is to develop a neighbourhood typology based on a series of demographic, socio-economic and environmental indicators as well as survey data in order to classify the urban neighbourhoods within the study area into clusters of similar type.

This process of dimensionality reduction and clustering of the high-dimensional attribute data serves to find matching pairs of treatment and control districts in order to enhance the validity of an observational study of a complex intervention across multiple sites. The clustering procedure to develop this neighbourhood typology thus has a double objective: on the one hand, the algorithm should reduce the within-cluster variance of the neighbourhood ecological variables, which may be correlated with the outcome variables and thus risk confounding inferences about program impact in a non-experimental research design. Put differently, neighbourhoods that resemble each other in terms of their demographic and socio-economic structure as well as the built environment must be grouped into clusters of similar type. On the other hand, the resulting neighbourhood typology should account for a maximum of the between-cluster variance in the outcome indicators prior to program implementation, i.e. the survey response patterns should be similar for residents of a given neighbourhood type across urban areas. In other words, the goal was to match the comparison areas not only on the observed covariates describing the neighbourhood context but on the outcome variables targeted by community policing such as fear of crime, neighbourhood disorder and satisfaction with police as well, i.e. the typology should capture a maximum of the *neighbourhood effects* of the different neighbourhood types. An observational study based on a neighbourhood typology that meets both these objectives allows an evaluator to dismiss a series of threats to the internal validity that otherwise beset a non-experimental research design.

B. The methodology of matching neighbourhood districts

The Home Office researchers who developed the BCU and CDRP families across England and Wales pre-selected 20 variables form the British census describing the demographic, socio-economic and the built environment characteristics of these areas based on their correlation with area levels of crime and disorder. As clustering algorithm they then employed k-means and self-organizing maps (SOM) in order to develop a typology of BCUs and CDRPs that minimized the variance in crime rates within a given family (Harper et al., 2002).

Outside criminology, two more recent studies used artificial neural networks and data mining procedures

to develop typologies of areal units of analysis. Li and Shanmuganathan (2007) used the SOM algorithm for a clustering of 90 demographic and socio-economic variables for a social area analysis to classify 163 census tracts in a small-sized city in western Japan. Spielman and Thill (2008) used self-organizing maps for a geo-demographic classification of 2,217 census tracts in New York City, using a dataset of 79 attributes from the U.S. Census.

The current study uses both unsupervised and supervised data mining algorithms to develop the neighbourhood typology. During the unsupervised learning phase, self-organizing maps are being used to cluster a high-dimensional data set in order to classify the neighbourhood areas into clusters of similar type (Skupin & Agarwal, 2008; Vesanto & Alhoniemi, 2000). During the following supervised learning phase, the random forests algorithm (Breiman, 2001) is used to select the most important features in order to develop a parsimonious model that makes a minimum of classification errors. In addition, the random forests algorithm serves as a gauge of the performance of the clustering algorithm overall as well as the predictive power of individual variables in the training data set.

As a final step, the resulting neighbourhood typology is to be visualized in a GIS as a map in the original geographic space, indicating the location of areas of a given type that are thus suited for matching and comparing during the subsequent impact evaluation.

III. DATA AND METHODOLOGY

The data analysed in the evaluation of community policing across Swiss urban areas come from three main sources: (a) official police crime statistics on area level crime rates, (b) the 1990 and 2000 Swiss population and housing census on the demographic composition and socio-economic status of the resident population as well as the structure of the built environment; and (c) the Swiss Crime Survey (SCS), a large-scale longitudinal criminal victimization survey on fear of crime, perceptions of neighbourhood disorder and popular attitudes towards the police. The SCS sampled the five urban areas under study repeatedly between 1998 and 2005. All data were measured at the level of postal ZIP code or administrative districts within the five urban areas, which are the smallest spatial unit of analysis for which all three data categories are available.

A. Unsupervised learning – Self-organizing maps

In geospatial data mining, unsupervised learning algorithms serve as analytical and modelling tools to discover patterns or structures in the data in order to classify study objects with (dis-)similar features in attribute or in geographic space (Kanevski et al., 2009). During the unsupervised learning phase, the current study uses self-organizing maps (SOM; Kohonen, 1990, 2001) as modelling tools to identify the underlying spatiotemporal patterns in the neighbourhood ecological data.

As a dimensionality reduction algorithm, the SOM is analogous to a discrete non-linear Principal Components Analysis (PCA). A PCA fits a hyper-plane into the data cloud that minimizes the distance to the original data points in order to replace the original variables by a smaller number of uncorrelated principal components. In the SOM algorithm, a network or *lattice* of artificial neurons is introduced into the input space instead of a hyper-plane. The segments of the SOM lattice are flexible and highly elastic and thus fit easily over curve-linear or unevenly distributed data during the training phase, a process which has been likened to covering the cloud of original data points with an *elastic fishing net* (Lee & Verleysen, 2007, 136).

After the training is completed, the artificial neurons or prototype vectors of the SOM lattice are aggregated by hierarchical agglomerative clustering (Skupin & Agarwal, 2008; Vesanto & Alhoniemi, 2000). As explained below, the resulting dendrogram is then cut at the level that produces the partitioning of the prototype vectors that best meets the twin objectives of the clustering algorithm. Finally, the original data points representing the urban neighbourhoods are assigned to the class of the nearest prototype vector, to which they have been attached during the training of the SOM lattice (Figure 1).

B. Supervised learning – Random forests

Supervised data mining algorithms seek structures or patterns in the data that explain or predict a priori information on outcomes or classifications (Kanevski et al., 2009). During supervised learning, the neighbourhood classification system that resulted from the unsupervised training is used as a priori information or labels of the training data. In this phase, the random forests algorithm (Breiman, 2001) serves to develop a classifier that aims to predict each neighbourhood's class based on the same explanatory variables that were first used during the unsupervised learning.

In the current context, supervised learning can be likened to non-linear multivariate logistic regression with the neighbourhood classification being the dependent variable. As a logistic regression equation predicts a value or class of the dependent variable, the random forests classifier contains a decision rule that assigns each neighbourhood area to the cluster to which it belongs based on its values on the explanatory or training variables.

The purpose of the supervised learning is on the one hand to assess the validity of the neighbourhood typology by analysing the overall classification error rate of the decision rule. As a by-product of the training of the classifier, the random forests algorithm produces a
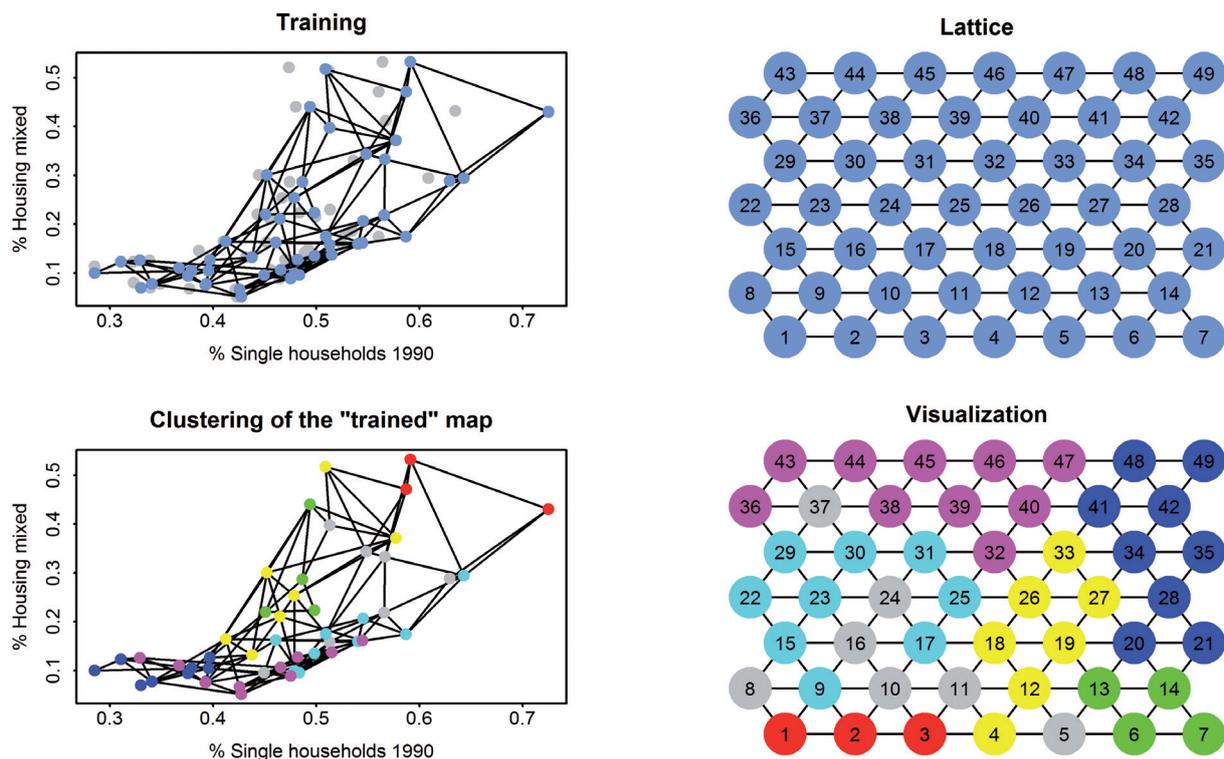
Enhancing the design of observational studies of community policing:
using geospatial data mining to design non-experimental program evaluations

63

Figure 1. The SOM algorithm visually explained. During training, the network of prototype vectors (blue) is iteratively fitted over the original data points (top left). The structure of the SOM lattice in projected 2-D output space (top right). Following training, the prototype vectors are clustered and labelled according to their topological position inside the lattice (bottom left). Visualization of the partitioning of the lattice in output space (bottom right). NOTE: Training and clustering occur in high-dimensional input space; for simplicity, only 2 variables are plotted here.

proximity measure that indicates the probability that any two neighbourhoods will be assigned to the same cluster based on their explanatory variables and thus how closely they resemble each other in the original input space (Breiman, 2001; Liaw & Wiener, 2002, 18f.). As a general rule, the greater this distance is for neighbourhoods belonging to different classes, the more distinct the neighbourhood clusters really are and the fewer classification errors the random forests classifier will make.

On the other hand, the random forests algorithm serves to weed out noisy variables with no or little predictive power in order to come up with a parsimonious model that makes a minimum of classification errors. As a second analytical tool the random forests algorithm computes a variable importance measure, which indicates the predictive value of each explanatory variable. Not unlike the *p*-value of a coefficient in multivariate regression, the variable importance measure can be used for feature selection to remove noisy or unimportant explanatory variables from the training data set (Breiman, 2001, 23f.; Genuer et al., 2010, 2226, 2229). This procedure is quasi analogous to a stepwise regression procedure which recursively removes non-significant variables from the regression model until only the most pertinent predictors are retained.

C. The neighbourhood clustering procedure

The methodological approach of the current study to develop a typology of Swiss urban neighbourhood areas uses both unsupervised and supervised data mining in an iterative procedure. During unsupervised learning, self-organizing maps are used to detect patterns in the neighbourhood ecological data. After the training of the lattice is completed, hierarchical agglomerative clustering (HAC) serves to merge the SOM prototype vectors (Skupin & Agarwal, 2008; Vesanto & Alhoniemi, 2000). The dendrogram resulting from HAC is then cut at the level that produces a partitioning of the neighbourhood areas that best satisfies the twin optimization criteria of the clustering procedure.

Once the optimum number of clusters has been determined, the trained SOM lattice is divided into as many segments. The SOM prototype vectors are classified depending on their position inside the SOM lattice. The original data points representing the urban districts take on the label of the SOM prototype vector to which they have been attached during training.

During the supervised learning phase, the resulting neighbourhood classification system is used to label the training data set. The random forests algorithm serves to develop a classification rule that assigns each urban

district to a neighbourhood cluster based on the explanatory variables describing the ecological context. The random forest classifier serves both to assess the overall quality of the neighbourhood typology as well as to select the explanatory variables with high predictive value.

The original training data set to characterize the neighbourhood ecology contained 89 variables, which can be regrouped into five distinct categories: official crime rates, population demography, socio-economic status, population heterogeneity and residential stability and the built environment. In a first step, the SOM-random forests algorithms described above were run separately on four of the five categories of ecological variables in order to select the key features of each. The selection criterion for this initial clustering procedure was fairly straightforward: keep as few explanatory variables as necessary without unduly increasing the classification error rate of the random forests decision rule.

The following section first gives a brief account of the preliminary clustering procedures of each variable category, before describing the procedure and outcomes of the final model of the key 24 variables in greater detail.

1. Crime rates
For the four cities for which neighbourhood-level crime data were available for analysis – Bern, Geneva, Lausanne and Zurich – local police statistics have been harmonized and the standardized neighbourhood-level crime rates calculated for eight different types of criminal offenses: homicides, assaults, burglaries, motor vehicle thefts, robberies, vandalism, extortion and threats.

As these standardized neighbourhood-level crime rates of the eight criminal infractions turned out be highly collinear, a principal components analysis was run to replace them by their component scores on the emerging principal components. Only the first two principal components had eigenvalues greater than one and were thus retained for the final neighbourhood clustering analysis, labelled as "Crime PC1" and "Crime PC2".

2. Population composition
The second set of ecological variables describes the demographic composition of the neighbourhood population. It includes the percentage of the total population by age group in 18 categories of 5-year intervals, from "0 to 4 years old" to "85 years old and above" as well as the percentage of both single and family households in the area. Demographic data from both the 1990 and 2000 census were included in the analysis. After the initial clustering procedure, six out of the 40 variables were retained for the final analysis: the percentage of children aged "5 to 9 years old" and "10 to 14 years old" from the 1990 census as well as the percentage

of "Single households" and "Families" from both the 1990 and 2000 census.

3. Socio-economic status
The third set of ecological variables describes an area's socio-economic status, measured as the percentage of the resident population by level of the highest educational achievement in seven categories ranging from "Mandatory schooling" to "University" (graduates) from the 1990 and 2000 census. The list of variables also included the percentages of the active working population residing in the area subdivided by eight professional categories of varying social prestige and remuneration, ranging from "Unskilled workers" to "Executives" from the 1990 and 2000 census. Of these 22 SES indicators, six were retained for the final analysis: the percentage of residents who had completed "Mandatory schooling" or an "Apprenticeship" from both the 1990 and 2000 census, as well as the percentage of university graduates and of residents employed in a "Middle management" position from the 2000 census.

4. Heterogeneity and residential stability
The fourth set of ecological variables included five variables measuring an area's degree of population heterogeneity and residential stability: the percentage of Swiss and foreign nationals among neighbourhood residents from the 1990 and 2000 census as well as the percentage of residents who in 2000 still lived at the same address as five years earlier. After the initial clustering, three variables were kept for the final analysis: the percentage of "Foreigners" among the resident population in 1990 and 2000 as well as the variable capturing the percentage of long-term residents in an area.

5. Built environment
The fifth set of ecological variables included in the analysis characterizes the built environment in a given area. A first set of nine variables measures the percentage of buildings of the total housing stock by height, ranging from "1 story" and "2 stories" up to "15 and more stories". A second set of eight variables captures the period of construction of the building units in the area, distinguishing eight different time periods from "before 1900", to "1900-1920" until the most recent period "1986-1990". Three more indicators were included in the analysis measuring the percentage of building units by their functional use, distinguishing between residential buildings, mixed housing complexes that include both apartments and offices or shops and non-residential buildings such as office complexes or commercial centres. All indicators describing the built environment were gathered from the 1990 census only. Seven variables proved important during the initial clustering run and were included in the final analysis: the three indicators on the functional use of buildings ("Residential", "Housing mixed" and "Non-housing"), three

indicators of building height ("2 stories", "6 stories" and "7-9 stories") as well as one variable indicating the construction period ("before 1900").

6. Final Clustering Procedure of the Key Variables

These initial clustering procedures served primarily to select the most important neighbourhood ecological features for each category of data. Based on their results, it was possible to reduce the number of variables in the training data set from 89 (87 ecological variables + 2 principal components of the crime data) to 24 (22+2) key variables, without an undue increase in the classification error rate of the resulting random forests classifier for each of the four data categories.

In a second phase, the identical clustering procedure - the SOM algorithm and hierarchical agglomerative clustering to classify the unlabelled data paired with the random forests algorithm to assess the quality of the clustering and to identify the most important features – was run on the 24 key variables retained for the final model. The diagnostic plots that resulted from this final clustering procedure are shown in Figure 2.

In order to determine the number of clusters that results in the optimum partitioning of the ZIP code or administrative districts, the dendrogram resulting from the hierarchical agglomerative clustering of the SOM prototype vectors of the final model is cut at different levels. For each possible number of neighbourhood clusters $k = 1,2…K$, statistical tests are calculated to determine the optimum partitioning. As has been stated previously, the resulting neighbourhood typology has to reconcile a double objective: the typology should cluster the neighbourhoods that are most similar in terms

of their ecological characteristics and simultaneously account for a significant share of the between-cluster variance in the survey measures of community policing impact, namely fear of crime.

Each of the two optimization criteria required a separate test. For the first problem, identifying the optimum number of clusters with regard to the neighbourhood ecological data, some kind of clustering validity index must be applied. In line with other studies (e.g. Vesanto & Alhoniemi, 2000), the current study used the Davies-Bouldin Index (DBI; Davies & Bouldin, 1979) to assess the quality of different partitions of the urban districts across the four cities.

For the second optimization problem, identifying the number of neighbourhood clusters that accounts for the biggest share of the between-cluster variance in the survey data, the Swiss Crime Survey respondents were regrouped according to each possible number of neighbourhood clusters $k = 1,2…K$. The actual statistical test was a χ2-independence test to determine whether response patterns of the pooled survey sample varied significantly by neighbourhood clusters on the fear of crime survey item.

Figure 2(e) plots the test statistic of both the DBI (left scale) and the *p*-value of the χ2-independence test (right scale) as a function of the number of neighbourhood clusters into which the neighbourhood areas are divided across the four urban areas. This chart reveals that there is no unique solution that meets both optimization criteria simultaneously. Regarding the neighbourhood ecological data, a partitioning into merely two clusters would be ideal, as the DBI is at its global minimum for $k = 2$. However, if the neighbourhood areas are divided into just two clusters, the χ2-independence test on the
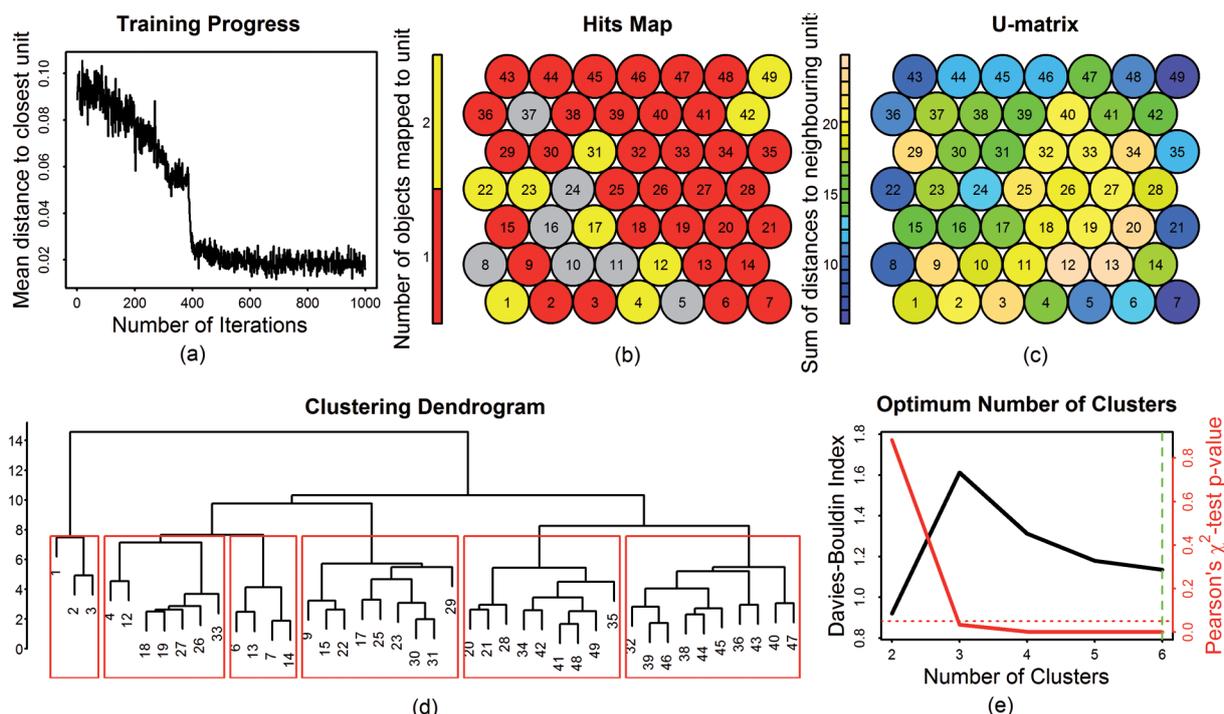

Figure 2. Training and clustering of the self-organizing map

survey data is not significant. A neighbourhood classification system that also accounts for a significant share of the between-cluster variance in the survey item requires at least three neighbourhood clusters. It turns out that the partitioning that best reconciles the twin optimization problems is six neighbourhood clusters: with $k = 6$, the χ2-independence test is significant and the test value of the DBI is at a local minimum.

Since the results of the SOM algorithm depend to some extent on the random initial values, the entire SOM training procedure and determination of the optimum number of clusters described above was replicated 50 times. From among the solutions with a minimum of four clusters, the number of neighbourhood clusters that the algorithm suggested most frequently was six. (A requirement of a minimum of clusters had to imposed since solutions with less than four clusters always pitted the urban centres against the rest of the urban areas, which failed to explain a significant share of the variance of the survey outcome measures). From among those replications that suggested $k = 6$ as the optimum number of clusters, the replication that resulted in the lowest value of the DBI was retained as the final clustering result and presented in Figure 2.

All computations were made using the R language for statistical computing (R Development Core Team 2012). The SOM and Random Forests algorithms were computed using the R packages 'kohonen' (Wehrens & Buydens, 2007) and 'randomForest' (Liaw & Wiener, 2002). Survey data were analyzed using the svymean and svytable functions of the package 'survey' (Lumley, 2010) and the Davies-Bouldin index was computed using the package 'clusterSim' (Walesiak & Dudek, 2012). Map illustrations and multi-panel plots were created using packages 'maptools' (Bivand et al., 2008) and 'lattice' (Sarkar, 2008), respectively.

IV. RESULTS

Once the optimum partitioning has been determined, the resulting neighbourhood typology can be visualized using different tools. Figure 2(d) plots the dendrogram of the hierarchical clustering of the SOM prototype vectors of the final model, cut at the optimum level $k = 6$. Since the leaves of the dendrogram represent the SOM prototype vectors, the result of the clustering procedure can also be visualized by means of the SOM lattice projected in 2D space with the different segments coloured according to the six neighbourhood types (Figure 3 top right panel). In addition, since each prototype vector represents one or more of the original urban areas, the ZIP code or administrative districts can be shaded according to the same colour scheme and the resulting neighbourhood typology visualized as maps of the five cities using GIS.

A. Geo-visualization of the neighbourhood typology

Figure 3 shows maps of the five urban areas. These maps display some striking parallels between the five cities. First of all, neighbourhood types 1 and 2 are the downtown areas located at or near the city centres. Neighbourhood types 3 and 4 form a first rim around the city centres, whereas neighbourhood types 5 and 6 are suburbs located on the outskirts of the five urban areas. Secondly, the five maps also reveal some striking differences, most notably between the French speaking cities of Lausanne and Geneva on the one hand and the Swiss German cities of Basel, Bern and Zurich on the other. Whereas in the French speaking areas, the suburban neighbourhoods are predominantly blue (type 5), the outskirts of the German speaking cities are either light blue or pink (type 4 and type 6; Kreis, 2012).

Moreover, the spatial pattern of the neighbourhood typology revealed some noteworthy characteristics both of the neighbourhood clusters themselves and the clustering algorithm used to identify them. First of all, in all five urban areas, the different neighbourhood types are neatly aligned on a centre-periphery axis. This pattern is all the more remarkable given that no geographic indicator was included as a variable in the training data used in the clustering algorithm. In other words, the variables included in the training data set on area-level crime rates, population composition, socio-economic status, heterogeneity and residential stability and the built environment display sufficient variation between the more central and the more peripheral areas for these neighbourhoods to cluster neatly into groups of similar type according to their geographic location (Kreis, 2012).

In a similar vein, the resulting spatial pattern of the neighbourhood typology also highlights the topology preserving quality of the SOM clustering algorithm. Topology preservation in a dimensionality-reduction algorithm means that two points that are close to each other in the high-dimensional original input space remain in close proximity of each other in the low-dimensional projected output space (Lee & Verleysen, 2007). This trait of the SOM algorithm is demonstrated neatly by the plot of the trained SOM lattice in the projected 2-D output space displayed in the top right panel of Figure 3. This chart shows the prototype vectors representing the urban centres or type 1 neighbourhoods at opposite ends of the SOM lattice from the suburban neighbourhoods of types 5 and 6, with the SOM prototype vectors representing the neighbourhood types 2, 3 and 4 lying in between. Again, it is noteworthy that this pattern resulted without imposing any geographic reference on the clustering algorithm. In other words, the trained SOM lattice accurately reflects the spatial logic following a centre-periphery axis inherent in the training data. This provides further evidence that the SOM algorithm is well suited for the clustering task at hand (Kreis, 2012).
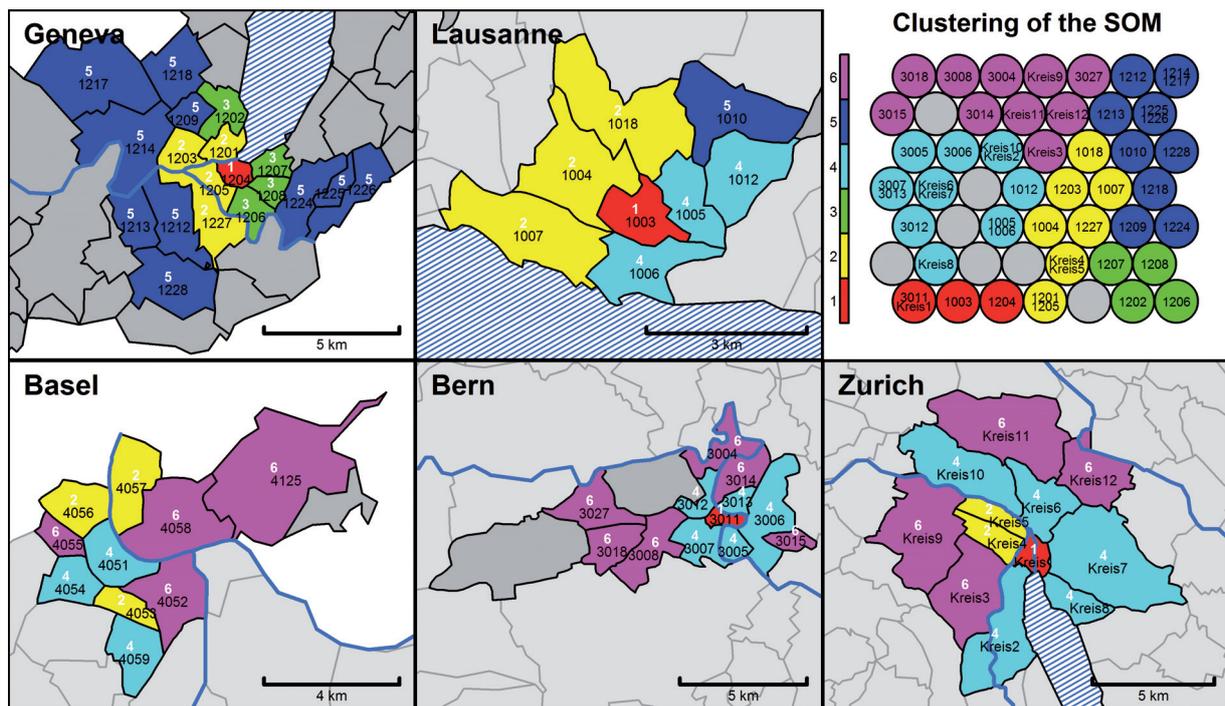
Figure 3. Visualization of the neighbourhood typology in geographic space. Maps of the five cities with neighbourhood areas coloured according to type. The top right panel displays the trained SOM lattice projected into 2-D output space with the segments coloured according to the best partitioning. Neighbourhoods in Basel were labelled based on the random forests classifier trained on the ecological data on the other four cities since for it no neighbourhood-level crime data were available.

B. Assessing the valdity of the neighbourhood typology for evaluation purposes

Once the clustering procedure is completed, the validity of the resulting neighbourhood typology for the planned evaluation of community policing still needs to be assessed. The first optimization criterion for the clustering algorithm was that the neighbourhood typology should minimize within-cluster variance in the neighbourhood contextual variables. In order to assess the clustering algorithm on this score, Figure 4 plots the attribute values of all the 51 neighbourhoods in the data sample. For each of the 24 variables retained in the final clustering algorithm to develop the neighbourhood typology, a panel in Figure 4 displays the original attribute values of the neighbourhoods as boxplots for each of the six neighbourhood types (Kreis, 2012).

The boxplots reveal that the clustering algorithm clearly reduced the within-cluster variance on each of the 24 key variables. To the extent that the boxplots by neighbourhood types are separated, Figure 4 highlights the underlying differences in the attribute between the different neighbourhoods. As was to be expected, not every variable is equally important in distinguishing the six neighbourhood types. Figure 4 thus also serves to identify the salient characteristics of each of the six neighbourhood types that most differentiate it from the other neighbourhood types (Kreis, 2012).

The second and arguably more important criterion of the validity of the neighbourhood typology is whether it accounts for a significant share of the between-cluster variance in the survey variables that serve as indicators of community policing impact. The survey item used as the optimization criterion in the clustering algorithm itself was the question on fear of crime, which thus varies significantly by neighbourhood type. It still has to be determined to what extent the neighbourhood typology also picks up some of the variance of the other survey outcome measures of community policing.

Figure 5 displays a chart for each of the six survey items that were used as the outcome indicators in the evaluation of community policing in Swiss urban areas: three items tapping into survey respondents' fear of crime, two indicators measuring physical and social disorder and one item capturing popular satisfaction with the police. For each of the survey items, the pooled SCS survey sample from all five urban areas was divided into six subgroups according to the neighbourhood type of a respondent's place of residence (type 1 to type 6). The six bar-plots represent the underlying contingency table of the neighbourhood cluster subgroup sample by the answering category to the survey item.

For each of the six survey items, the panels indicate the survey question with the corresponding answering categories. The bar-plots display the percentage of survey respondents by answering category. The single
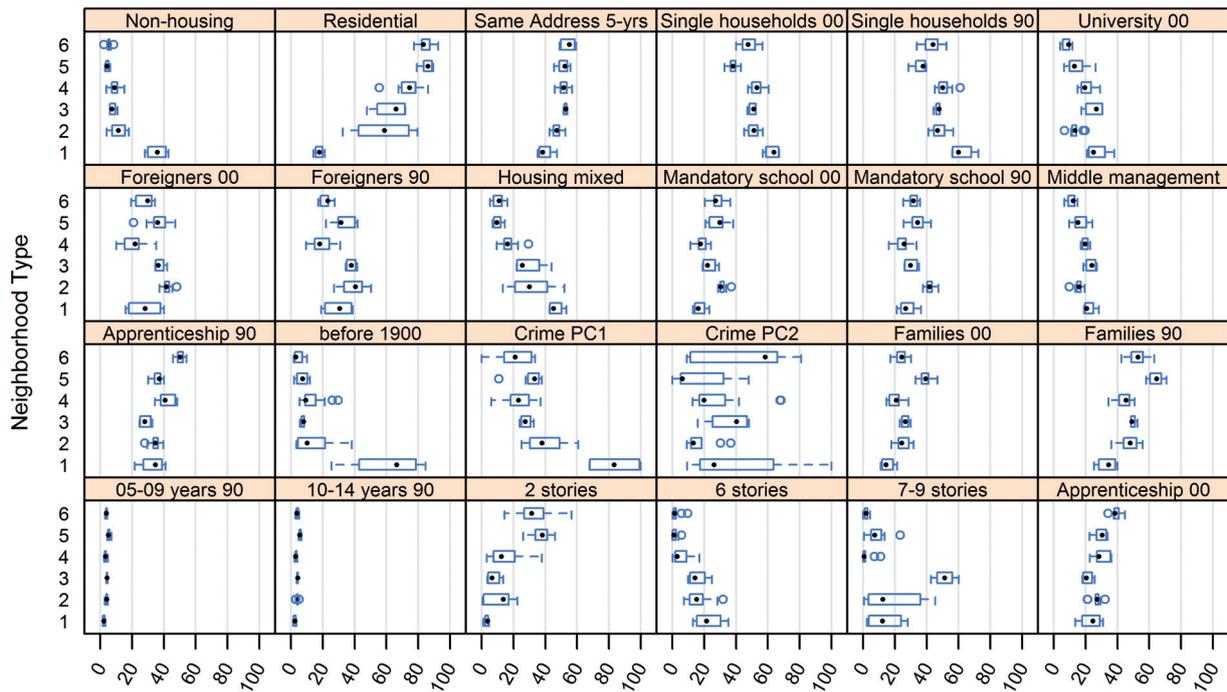
Figure 4. Defining characteristics of the neighbourhood typology. Original values of the 24 variables retained in the final clustering analysis to describe the neighbourhood ecology. All variables are percentages except for the crime principal components scores ("Crime PC1" and "Crime PC2"), whose true range was linearly transformed to a 0-100 scale.
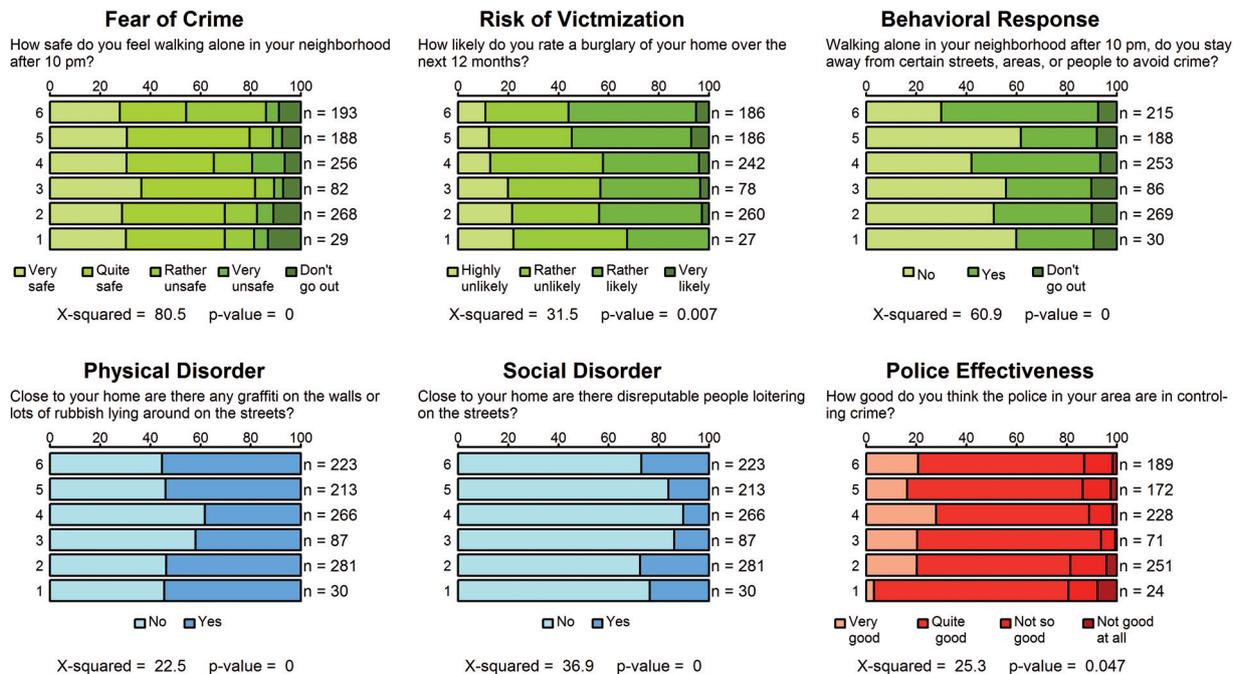

Figure 5. Survey response patterns by neighbourhood type. Percentage of respondents by answer category for the six survey items used to assess the impact of community policing on neighbourhood residents. Survey respondents were grouped together across cities by neighbourhood type, excluding respondents from Basel. The χ2-independence test statistics were calculated using Monte Carlo simulations. The total survey sample was weighted, stratified at the neighbourhood level, to correct for sampling bias in the age and gender distribution.

Enhancing the design of observational studies of community policing:
using geospatial data mining to design non-experimental program evaluations

69

digits on the left of each bar-plot indicate the type of neighbourhood cluster, whereas the numbers on the right indicate the size of the subgroup sample of each neighbourhood cluster. The survey data were weighted as a stratified random sample at the ZIP code level to correct for sampling bias in the age and gender distribution to make the neighbourhood-level sub-samples representative of the local resident population. At the bottom of each panel are the test statistics of the $\chi2$-independence test that was run to determine whether survey response patterns differ significantly by neighbourhood type. These test statistics were calculated by means of a Monte Carlo simulation to circumvent the problem of contingency table cells with an expected frequency of below five. As a matter of fact, the absolute number of survey respondents per answering category is at times rather low, especially for neighbourhood type 1, i.e. the urban centres. Monte Carlo simulations are implemented as a standard option in the chisq.test function in R and were computed on the basis of 2000 replications, which is the standard value R proposes for such simulations.

The contingency tables behind the boxplot charts and corresponding $\chi2$-independence tests unearthed some very interesting spatial trends in the survey response patterns that deserve a closer inspection. As a matter of fact, the boxplots of the six different neighbourhood types are arranged in a spatial order inside each chart, with the more centrally located neighbourhood clusters (types 1, 2 and 3) being at the bottom and the more peripheral neighbourhood (types 4, 5 and 6) placed on top.

The top left panel shows the bar-plots of the survey item of fear of crime measured as the feeling of safety on a nightly stroll through one's own neighbourhood. As this chart indicates, the percentage of respondents who feel "very safe" or "quite safe" is generally higher in the more centrally located neighbourhoods than in the outskirts. The only notable exceptions to this spatial trend are neighbourhood clusters type 3 and 5, which are predominantly neighbourhoods located in the Geneva urban area (Kreis, 2012).

This general tendency of fear of crime to be more prevalent in the outer areas than in the urban centres is still more pronounced for the survey item asking about the perceived risk of a burglary of one's home. The percentage of respondents who rate the chances of a burglary of their home over the next twelve months as "likely" or "very likely" goes up systematically from neighbourhood type 1 to 6, i.e. as one moves out from the city centres to the outskirts of the five urban areas in the real world (Kreis, 2012).

For burglary, the spatial pattern of the perceived risk can be compared to actual victimization risk as captured by the official police crime statistics. An earlier study that mapped the neighbourhood-level burglary rates for four of the five urban areas under study here found that the relative risk of a residential burglary was more acute in the urban centres and tends to decrease the further one moves away from the downtown areas. The spatial pattern of the survey responses, however, moves in exactly opposite direction. In other words, neighbourhood residents collectively are rather poor at evaluating victimization risk: the risk of a burglary is underrated in the city centre and overrated in the peripheries (Kreis, 2012).

The general spatial trend of fear of crime to increase from centre to periphery also applies to the response pattern for the third indicator, actual behavioural changes to avoid crime. The percentage of respondents who resort to behavioural changes to avoid crime in their neighbourhood steadily increases from the centrally located neighbourhood clusters towards the urban peripheries. The only outlier in this general spatial trend is again neighbourhood cluster type 5, which are the suburban areas of Geneva and Lausanne (Kreis, 2012).

The first two panels of the bottom row of Figure 5 show the bar-plots for the survey item measuring physical and social disorder. For the disorder items, the spatial pattern is no longer a more or less linear trend that increases from the centre to the periphery as with fear of crime. The percentage of respondents who spotted signs of physical or social disorder, or both, is higher for the more central neighbourhood clusters type 1 and 2 as well as the peripheral areas type 5 and 6, but slightly lower for the in-between areas of type 3 and 4 (Kreis, 2012).

The sixth panel on the survey item measuring popular satisfaction with the police reveals a spatial pattern that is more in line with actual victimization risk. Asked whether local police were doing a satisfactory job in crime control, the percentage of respondents who rated the police as doing a "very good" or "quite good" job increases from centre to periphery, or from neighbourhood clusters type 1 to 6. This is in line with the spatial pattern detected in an earlier analysis of the neighbourhood-level burglary victimization rates in Swiss urban areas (Kreis, 2012).

The second criterion of the clustering algorithm to develop the neighbourhood typology – that it accounts for a significant share of the between cluster variance in the survey outcome measures – has thus been met. However, before the current neighbourhood typology may be used as intelligence basis to select matched treatment and control areas across urban areas to study the impact of different community policing strategies, a final check is still in order. This test must assess whether the current typology does indeed account for most of the variance in the outcome variables between residents of different neighbourhood types or if there is a significant amount of variance left at the higher aggregate city or regional levels.

In order to test this proposition, a second series of $\chi2$-independence tests is run, this time to evaluate whether

the response patterns of the survey respondents of a given type of neighbourhood cluster vary significantly between individual cities. Table 1 displays all *p*-values of these tests for all six survey outcome indicators for each of the six neighbourhood types. The results are encouraging: the *p*-value of a majority of these χ2-independence test is not significant, suggesting that the response patterns of survey respondents residing in the same type of neighbourhood do not differ significantly between cities and that the null hypothesis of basic independence cannot be rejected (which is what the author was hoping for). However, several χ2-independence tests have a *p*-value that is significant. As a matter of fact, for all five neighbourhood clusters, which are present in more than one urban area, is the *p*-value below the conventional 0.05 significance level on at least one occasion. This implies that for those neighbourhood types and those survey items, city-level factors still impinge on survey response patterns (Kreis, 2012).

V. DISCUSSION

The current research employed both unsupervised and supervised data mining algorithms to develop a typology of neighbourhoods in order to group neighbourhood districts across the major Swiss urban areas into clusters of similar type. Since the objective of the clustering algorithm was to match suitable *treatment* and *control* areas across the five urban areas in order to enhance the internal validity of an observational study of community policing, the twin optimization criteria for the clustering algorithm were clear: on the one hand, the neighbourhood typology should minimize the between-cluster similarity in the contextual variables, which are potentially correlated with the outcome measures, and thus reduce the risk that these neighbourhood covariates confound any inferences about the impact of the treatment. On the other hand, the neighbourhood typology should account for a significant share of the between-cluster variance in the outcome measures used to evaluate community policing. This was meant to ensure that residents of the same neighbourhood type in different urban areas collectively expressed similar views prior to the onset of treatment and later observed differences in opinion are not due to pre-existing conditions at the onset of community policing implementa-

tion (Kreis, 2012).

As the diagnostic plots of the previous section revealed, these twin optimization criteria have largely been met: not only did the neighbourhood typology reduce between-cluster variance in the 24 key variables describing the neighbourhood ecological context. It also accounted for a significant share of the between-cluster variance in the survey response patterns that served as outcome indicators in the evaluation of community policing.

The neighbourhood classification system developed for the community policing evaluation thus goes a long way to reconcile these double optimization criteria notwithstanding the fact that the algorithm did not result in a single optimum number of clusters. Indeed, one conclusion from the clustering procedure was that no such single optimum number of neighbourhood clusters exists. The number of neighbourhood categories that has led to the most clear-cut separation of the clusters on the ecological data is smaller than the number of neighbourhood clusters that accounts for the biggest amount of variance in the outcome survey measures. This implies that in the present case the individual clusters of neighbourhoods of similar type are less distinct and tend to overlap in the original input space of the ecological data, so that individual neighbourhoods could be classified either way (Kreis, 2012). The apparent lack of a single optimum number of neighbourhood clusters compounds the variability of the results inherent in the SOM clustering algorithm. The shape to which the malleable SOM lattice converges during training depends to some extent on the randomly chosen initial values of the weights of the prototype vectors. The standard remedy to handle this aspect of the SOM algorithm is to replicate the clustering procedure multiple times and to compare the results of the individual runs before reaching any conclusions. For the current study, the complete clustering algorithm was replicated 50 times and the solution retained that best met the double optimization criteria set at the outset of the clustering procedure.

By contrast, the MC simulated χ2-independence tests as well as the random forests algorithm that were both based on the outcome of the SOM algorithm produced very stable results. Random forests were employed during the supervised learning phase in order to identify the key ecological variables among the indicators used

	Fear of Crime	Risk of Victimization	Behavioural Response	Physical Disorder	Social Disorder	Police Effectiveness
1	0.007	0.824	0.006	0.019	0.001	0.857
2	0.000	0.000	0.494	0.395	0.195	0.164
3						
4	0.001	0.130	0.137	0.000	0.007	0.034
5	0.817	0.107	0.130	0.037	0.262	0.367
6	0.301	0.000	0.590	0.734	0.060	0.052

Table 1. Survey response patterns by neighbourhood cluster. *p*-values of the χ2-independence tests of the indicators of community policing impact by city, computed separately for each of the six neighbourhood clusters (a value of *p* < 0.05 indicates that response patterns within a given neighbourhood cluster still vary significantly by city).

to describe the neighbourhood context. By selecting only the most important indicators it was possible to build a parsimonious final model with just 24 explanatory variables (out of the original 89) without unduly increasing the classification error rate of the model.

A second limitation of the neighbourhood typology is that it could not account for all or most of the variance in the survey measures used as outcome indicators at the higher aggregate levels of analysis. If survey respondents are grouped by individual neighbourhood clusters, there remains at times significant within-cluster variance in the response patterns at the city-level. This implies that city-level factors still influence survey response patterns, which risks undermining comparisons between neighbourhood residents across urban areas even within a given neighbourhood type. In an observational study design that compares the impact of the program between neighbourhoods of a similar type across urban areas, the current neighbourhood typology thus manages to reduce the threats to internal validity of selection and regression to the mean but cannot rule them out completely (Kreis, 2012).

VI. CONCLUSION

The current research employed geospatial data mining algorithms to classify Swiss urban neighbourhoods into clusters of similar type in order to find matching treatment and control areas for the evaluation of area-based crime prevention programs such as community policing. The clustering procedure made it possible to take high-dimensional data on the demographic and socioeconomic composition as well as the built environment of urban neighbourhoods into account. Not only were these data shown to impact survey response patterns, they may exert an influence on how neighbourhood residents perceive different community policing strategies as well. This approach thus attempted to blunt some of the criticism levelled against recent criminological research on crime prevention of being overly concerned with the question of *what works?* while neglecting the influence of contextual and environmental factors on the success of such initiatives (Williamson et al., 2006, 199f.).

Despite some shortcomings previously discussed, the resulting neighbourhood typology succeeded in achieving a considerable degree of within-cluster homogeneity regarding the ecological data while capturing a significant share of the between cluster variance in the survey items. The individual neighbourhoods within each cluster thus not only share similar ecological characteristics, which may confound inferences about the impact of treatment, but also display similar levels on the outcome measures prior to program implementation, which community policing is trying to influence. The neighbourhood typology thus helps to diminish some of the threats to internal validity of an observa-

tional research design and first make possible valid comparative analysis of the impact of area-based crime prevention programs across urban areas.

Besides diminishing the threats to internal validity of an observational study design, the neighbourhood typology unearthed some peculiar facts that are also interesting from a policy perspective. First of all, the typology suggests that neighbourhood residents collectively are not very perceptive when it comes to assessing victimization risk. Whereas actual (burglary) risk is highest in the city centres and tends to diminish towards the boundaries of the urban areas, the spatial pattern for perceived risk is exactly the reverse. The study thus produced evidence that also in Swiss urban areas there are neighbourhoods that show signs of the phenomenon that has become known as the *reassurance gap* (Tuffin et al., 2006), meaning a popular perception of an increase in crime when actual rates are low or falling. This is not to say that neighbourhood residents' perception are completely off the mark, however: when asked about whether the local police are doing a good job in controlling crime in the area, popular approval rates do co-vary spatially with actual levels of victimization (Kreis, 2012).

The second extra bit of information the neighbourhood classification system generated was to highlight the key characteristics of each neighbourhood type from the high-dimensional training data set of ecological indicators. The random forests algorithm proved highly informative in this regard, identifying the key characteristics of each of the six neighbourhood clusters. Needless to say that such information can be used to tailor policy interventions to the specific needs of these different areas.

ACKNOWLEDGEMENTS

REFERENCES

Ashby, D.I. and Longley, P.A. (2005). Geocomputation, geodemographics and resource allocation for local policing. *Transactions in GIS*, *9* (1), 53-72.

Bivand, R.S., Pebesma, E.J., and Gómez-Rubio, V. (2008). *Applied Spatial Data Analysis with R*. New York: Springer.

Breiman, Leo. (2001). "Random forests." *Machine Learning*, Vol. 45, No. 1, pp. 5-32.

Cook, T.D. and Campbell, D.T. (1979). *Quasi-experi-*

mentation: Design and Analysis Issues for Field Settings. Chicago: Rand McNally College Publishing Company.

Davies, D.L. and Bouldin, D.W. (1979). "A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *1* (2), 224-227.

Farrington, D.P. (2003). Methodological quality standards for evaluation research. *The ANNALS of the American Academy of Political and Social Science*, 587 (1), 49-68.

Genuer, R., Poggi, J.-M. and Tuleau-Malot, C. (2010). Variable selection using random forests. *Pattern Recognition Letters*, *31* (14), 2225-2236.

Harper, G., Williamson I., Clarke, G. and See, L. (2002). *Family Origins: Developing Groups of Crime and Disorder Reduction Partnerships and Police Basic Command Units for Comparative Purposes*. London: Home Office.

Kanevski, M., Pozdnoukhov, A. and Vadim T. (2009). *Machine Learning for Spatial Environmental Data: Theory, Applications and Software*. Lausanne: EPFL Press.

Kohonen, T. (1990). The self-organizing map. *Proceedings of the IEEE*, 78 (9), 1464-1480.

Kohonen, T. (2001). *Self-Organizing Maps* (3rd ed.). Berlin: Springer.

Kreis, C. (2009). *The Spatio-Temporal Patterns of Fear of Crime in Major Swiss Cities: A quantitative analysis of the spatial distribution of fear of crime, disorder, and public satisfaction with the police among neighborhood residents in Switzerland's five metropolitan areas between 1987 and 2005* (Master thesis in études avancées en urbanisme durable, University of Lausanne, Lausanne).

Kreis, C. (2012). *Community Policing in Switzerland's Major Urban Areas: An Observational Study of the Implementation and Impact Using Geospatial Data Mining* (Ph.D. thesis, University of Lausanne, Lausanne).

Lee, John A. and Michel Verleysen. (2007). *Nonlinear Dimensionality Reduction*. New York: Springer.

Li, Y. and Shanmuganathan, S. (2007). *Social Area Analysis Using SOM and GIS: A Preliminary Research* (RCAPS Working Paper, Vol. 07-3). Beppu City, Japan: Ritsumeikan Asia Pacific University, Ritsumeikan Center for Asia Pacific Studies.

Liaw, A. and Wiener, M. (2002). Classification and regression by randomForest. *R News*, *2* (3), 18-22.

Lumley, T. (2010). *Complex Surveys: A Guide to Analysis Using R*. Hoboken: Wiley.

R Core Team (2012). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rosenbaum, P.R. (2010). *Design of Observational Studies*. New York: Springer.

Sarkar, D. (2008). *Lattice: Multivariate Data Visualization with R*. New York: Springer.

Shadish, William R., Cook, Thomas D., & Campbell, Donald T. 2002. *Experimental and Quasiexperimental Designs for Generalized Causal Inference*. Boston: Houghton Mifflin.

Sheldon, G., Hall, R., Brunsdon, C., Charlton, Alvanides, M.S. and Mostratos, N. (2002). *Maintaining basic command unit and crime and disorder partnership families for comparative purposes* (RDS Online Report).

Sherman, L.W., Farrington, D.P., Welsh, B.C. and MacKenzie, D.L. (Eds.). (2002). *Evidence-Based Crime Prevention*. London: Routledge.

Skupin, A. and Agarwal, P. (2008). Introduction: What is a self-organizing map? In Agarwal, P. and Skupin, A. (Eds.), *Self-Organising Maps: Applications in Geographic Information Science* (p.1-20). Chichester: Wiley.

Spielman, Seth E. and Jean-Claude Thill. (2008). "Social area analysis, data mining, and GIS." *Computers, Environment and Urban Systems*, Vol. 32, No. 2, pp. 110-122.

Tuffin, R., Morris, J. and Poole, A. (2006). *An Evaluation of the Impact of the National Reassurance Policing Programme* (Home Office Research Study, Vol. 296). London: Home Office.

Vesanto, J. and Alhoniemi, E. (2000). Clustering of the self-organizing map. *IEEE Transactions on Neural Networks*, *11* (3), 586 - 600.

Walesiak, M., and Dudek, A. (2012). *clusterSim: Searching for optimal clustering procedure for a data set*. R package version 0.41-8.

Weisburd, D. and Eck, J.E. (2004). What can police do to reduce crime, disorder, and fear? *The ANNALS of the American Academy of Political and Social Science*, *593* (1), 42-65.

Wehrens, R., and Buydens, L. (2007). Self- and super-organizing maps in R: the kohonen package. *Journal of Statistical Software*, 21(5), 1–19.

Welsh, B.C. and Hoshi, A. (2002). Communities and crime prevention. In Sherman, L.W., Farrington, D.P., Welsh, B.C. and MacKenzie, D.L. (Eds.), *Evidence-Based Crime Prevention* (p 165-197). London: Routledge.

Williamson, T., Ashby, A.I. and Webber, R. (2006). Classifying neighbourhoods for reassurance policing. *Policing and Society*, *16* (2), 189-218.

Coordonnées de l'auteur :

Christian Kreis
Nederlands Studiecentrum Criminaliteit
en Rechtshandhaving
De Boelelaan 1077a
1081 HV Amsterdam
ckreis@nscr.nl