


Neighbourhood change and spatial inequalities in Cape Town

Christopher D. Lloyd¹  | Saad Bhatti² | David McLennan³ | Michael Noble³ | Gerbrand Mans⁴

¹Geography, School of Natural and Built Environment, Queen's University Belfast, Belfast, UK

²School of Geography and Environmental Sciences, Ulster University, Coleraine, UK

³Southern African Social Policy Research Insights, Hove, UK

⁴Council for Scientific and Industrial Research, Pretoria, South Africa

Correspondence

Christopher D. Lloyd, Geography, School of Natural and Built Environment, Queen's University Belfast, Belfast, UK.
Email: c.lloyd@qub.ac.uk

Funding information

Research Councils UK > Economic and Social Research Council [Grant number: ES/N014022/1]

Abstract

Capturing the dynamics of population change in urban areas necessitates access to geographically fine-grained and temporally consistent data for several time points. Such data are generally not available and they must be created using standard population data which cannot usually be compared across time periods. In this paper, the focus is on changing spatial inequalities in Cape Town, South Africa. This paper details an approach to generation of gridded population counts (250 m by 250 m) for two census years – 2001 and 2011. Census data for Small Area Layers (SALs), Spot Building Count (SBC) data, and Open Street Map (OSM) landuse data were used to construct a grid of populated cells to which population counts are then reallocated. The reallocation of population counts from SALs to grid cells was undertaken using area-to-point kriging – an approach which is informed by the spatial variation in the population groups of interest as measured using the variogram. A case study based on population grids of unemployment rates shows how the grids can be used to chart changes and also to measure spatial inequalities across the city at two time points. The advantages of grids for capturing fine-scale complexities and correctly accounting for physical separation between communities are demonstrated and the results show, while the broad patterns of inequality are consistent across time, there are pronounced increases in inequalities in some neighbourhoods. These areas – and what leads some areas to fall further behind – should be the subject of attention by policy-makers.

KEYWORDS

census, grids, inequalities, scale, South Africa

1 | INTRODUCTION

The demise of Apartheid in South Africa meant the removal of racially discriminatory restrictions on population movement, which accelerated migration from the former homelands to the major cities, particularly in Gauteng and the Western Cape. Cape Town has experienced substantial population growth over the last three decades as a result of rural–urban migration and natural change. The pace, nature and form of this growth poses serious challenges in terms of its impact on inequality because it tends to reinforce existing concentrations of poverty and exclusion, and reproduce established social and spatial divisions. Constrained access to urban land, housing and public services means that the poor are often forced to settle in marginalised areas

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

The information, practices and views in this article are those of the author(s) and do not necessarily reflect the opinion of the Royal Geographical Society (with IBG).

© 2021 The Authors. *The Geographical Journal* published by John Wiley & Sons Ltd on behalf of Royal Geographical Society (with the Institute of British Geographers)

in and around the existing townships, therefore inequalities by income and socio-economic status generally may be exacerbated. Income inequalities can be captured using the Gini coefficient (see Keeley, 2015), which takes values between zero and one. Increasingly large values of the Gini coefficient indicate that income is concentrated amongst smaller and smaller numbers of people. Cape Town's Gini coefficient (for income) was very high at 0.57 in 2010,¹ and on a par with other municipalities in the Western Cape (Western Cape Government Provincial Treasury, 2012). This is transcribed into a starkly polarised spatial form with marked differences in deprivation levels between many spatially proximate (neighbouring) areas; that is, there are *spatial* inequalities. The majority of quantitative research concerning inequality in South Africa has utilised the “classical” measures of (income) inequality, including the Gini coefficient, expressed at national or occasionally provincial level. Whilst such measures are powerful tools for assessing change in inequality over time within South Africa, or for national comparisons, they say little about people's day-to-day “lived experience” of inequality. In order to understand how inequality affects people's lives (and how this in turn affects the functioning of local communities and wider city regions), it is necessary to understand how people actually experience inequality (including where and when) as they go about their daily lives. This paper focuses on the analysis of changing deprivation (specifically unemployment) and spatial inequalities in Cape Town between 2001 and 2011.

The research presented here was undertaken as part of the Economic and Social Research Council (ESRC) project “Changing Socio-Spatial Inequalities: Population change and the lived experience of inequality in urban South Africa” (InequalitiesSA). The project seeks to build an enhanced understanding of spatial inequalities in South Africa. This paper describes the first stages of this work, which entailed profiling the populations of small areas, mapping change in populations over these areas, and measuring spatial inequalities between neighbourhoods. To assess the nature of changes in spatial inequalities (and population change more generally) it is necessary to have data which are closely comparable – with common variable definitions and common geographical areas. The boundaries used to report census data tend to change through time as population numbers and densities alter, and so standard census outputs cannot usually be compared directly and it is generally not possible to show how one small area changed between one time point (e.g., 2001) and another (e.g., 2011). One solution is to reallocate the census counts from their original geographical units to a single comparable set of geographical zones which are the same for all time points. The paper first outlines some approaches to standardising data by creating population surfaces (grids) so that small areas can be compared across time. Next, existing grid-based population datasets with coverage of South Africa are reviewed. The focus then moves on to the study area – the city of Cape Town – and available data sources. The generation of population surfaces in the case of Cape Town is then detailed, followed by some preliminary analyses of spatial inequalities using the grids generated. The paper concludes by considering some key next steps and by highlighting some ways in which the concepts and methods presented are of relevance in other contexts.

2 | ANALYSIS OF CHANGE THROUGH TIME

Exploration of changes in small areas over time using census data is hampered by changing variables and definitions, and by changing zonal systems. Taking the case of South Africa, the small areas used to report population counts from the censuses of 1996, 2001 and 2011 are not the same. Thus, it is not possible to map population changes over small areas for these three censuses. Incompatible zonal systems can be matched using an *areal interpolation* procedure. The most straightforward approach to change of support (converting from a set of source zones to another set of zones or points) is to use a simple geographical information system (GIS) overlay procedure. This entails overlaying the *source* zones and the zones to which counts will be reallocated (the *target* zones). Taking the example of population counts, the areas of overlap are then used to compute how much of the population of a source zone will be assigned to a target zone. For example, if a target zone (e.g., 2011 zones) contains 25% of a given source zone (e.g., 1996 zones) then it will receive 25% of the source zone population, plus the relevant proportion of the population of any other overlapping source zones. This simple approach can be adapted so that additional information – such as land use type – can be used to refine allocations between zones. In a simple case, areas of water may receive no people, while sparsely populated rural areas will receive less people than similarly sized but more densely populated more urban areas.

Here, as detailed above, the objective was to create gridded estimates of census variables to allow for analyses of changes between 2001 and 2011. Grids can be created from irregular source zones, such as Small Area Layers (SALs; detailed below), using areal weighting, but a wide range of alternatives exist (see Lloyd, 2014 for a summary). Lloyd et al. (2017) used a combination of postcode densities and areal weighting followed by a smoothing process to create gridded counts of census variables for the censuses of 1971–2011 inclusive in Britain. Kernel smoothing approaches (e.g., see Martin, 1989, 1996) can also be used to distribute population counts from existing geographies to a new gridded geography. Another alternative is to use a geostatistical (kriging-based) approach to changing data scales (see Goovaerts, 2008). With a geostatistical approach, the reallocation of counts is informed by the spatial structure of the variable of interest – as captured using the variogram. For example, if

similar values cluster over small areas, then the approach should be different to a case where values are very similar over a larger area. To clarify, kriging (and other smoothing approaches) assumes that a grid square located within a high-unemployment SAL but on the boundary with a lower-unemployment SAL is likely to have a lower unemployment rate than a grid square within the same high-unemployment SAL but not close to the boundary of the lower-unemployment SAL. In effect, to compensate for the grid square near the SAL boundary having a lower than average unemployment rate for the SAL, kriging increases the unemployment rate for the grid square which is not close to the boundary, so that the grid squares taken together have the same average unemployment rate as the SAL as a whole. It is important to note that this “spatial smoothing” process will not necessarily apply everywhere across Cape Town as, in some cases, it is actually the periphery of the poor neighbourhoods that see the growth of informal settlements and, in such cases, spatial smoothing might reduce spatial contrasts which we might seek to retain. However, inspection of the results presented later shows little evidence of this having any notable impact. Figure 1 gives an example, showing how estimates within source zones (irregular boundaries) are not the same, and estimates at boundaries of neighbouring zones are more alike. With a conventional areal interpolation, all values attached to grid cell centres within a given source zone would be the same.

2.1 | Geostatistical approaches

The variogram (sometimes referred to as the semivariogram) is a measure of the degree to which values differ according to how far apart they are (i.e., spatial distance between data points). Observations (e.g., census area centroids) are separated by a given distance and direction which is termed the spatial lag. As an example, two observations may be separated by 2 km and one of these observation locations may be located directly north of the other observation location. The variogram is estimated



FIGURE 1 Kriging estimates of unemployment rates for a sample area

by computing the squared differences between all of the paired observations, and half the average value is obtained for all observations separated by a given spatial lag. Note that a specified lag tolerance (e.g., 2 km ± 1 km) is used where the observations are not located on a regular grid). The experimental variogram $\hat{\gamma}(\mathbf{h})$ for spatial lag \mathbf{h} is computed with:

$$\hat{\gamma}(\mathbf{h}) = \frac{1}{2p(\mathbf{h})} \sum_{i=1}^{p(\mathbf{h})} \{z(\mathbf{x}_i) - z(\mathbf{x}_i + \mathbf{h})\}^2 \tag{1}$$

where $z(\mathbf{x}_i)$ is the observation (e.g., unemployment rate) at location \mathbf{x}_i and $p(\mathbf{h})$ is the number of paired observations separated by the lag distance \mathbf{h} . A mathematical model can be fitted to the experimental variogram, most commonly using a fitting procedure such as weighted least squares. Models are usually selected from a set of “authorised” models (Webster & Oliver, 2007) and these comprise bounded and unbounded models. Bounded models level out as they reach a particular lag (that is, they have a sill (a finite variance)), while unbounded models do not reach an upper bound. The components of a bounded variogram model are shown in Figure 2. The nugget effect c_0 represents unresolved variation; this can include spatial variation at a distance smaller than the sample spacing, and also measurement error. Spatially correlated variation is captured by the structured component, c . The sill (sill variance) comprises the nugget effect plus the structured component ($c_0 + c$; the *a priori* variance). The spatial scale (or frequency) of spatial variation is represented by the range, a . As an example, if unemployment rates differ markedly over short distances, then the rates have a high frequency of spatial variation (or a short range, a). In contrast, if the rates are very similar over large distances (i.e., values vary regionally but not locally), then the rates have a low frequency of spatial variation. The structured component captures the *magnitude* of variation, while the range represents the spatial *scale* of variation. The variogram is a function of the data support – the variogram estimated from data over (for example) census areas is termed the areal variogram. The derivation of a point support variogram from the areal variogram is outlined below.

The most widely used variant of kriging is ordinary kriging (OK). OK predictions are weighted averages of the n nearest neighbours of the prediction location. The weights are determined using the coefficients of a model fitted to the variogram (or another function such as the covariance function).

2.2 | The change of support problem

In social science contexts, data are often available for zones rather than points. While individual or household level data are available in some contexts, these are usually provided without detailed spatial information, and spatially aggregated data usually offer the only means of exploring detailed spatial patterns. The data support, v , is defined as the geometrical size, shape, and orientation of the units associated with the measurements (Atkinson & Tate, 2000). Thus, making predictions from areas to points corresponds to a change of support. Geostatistics offers the means to (1) explore how the spatial structure of a variable




FIGURE 2 Transitive (bounded) variogram model

changes with change of support, and (2) change the support by interpolation to an alternative zonal system or to a quasi-point support (Schabenberger & Gotway, 2005). For many applications, the variogram defined on a point support cannot be obtained and only values over a positive support (area) may be available. The variogram of aggregated data is termed the regularised or areal variogram (see Goovaerts, 2008).

Variogram deconvolution

If the point support variogram is available then the variogram can be estimated for any support. Using a variogram deconvolution procedure, the point support variogram can be estimated using the areal variogram. An iterative procedure was implemented by Atkinson and Curran (1995) to derive the point support variogram from the variogram estimated from data on regular grids. In population geography, variogram deconvolution for irregular supports (for example, census or administrative zones), rather than regular cells, is likely to be more useful. A variogram deconvolution approach for data on irregular supports is detailed by Goovaerts (2008). As for approaches used with regular supports, the objective of the method is to minimise the difference between the regularised variogram, which is derived from the punctual (deconvolved) variogram, and the variogram estimated from the areal data. Goovaerts summarises a 10-step procedure for deconvolution of the regularised variogram. This method is implemented in the SpaceStat software (see <http://www.biomedware.com/>). An alternative approach is available through the R package *rtop* (Skøien et al., 2014).

Area-to-point kriging

Given the deconvolution procedures outlined above, it is possible to make predictions at point locations using data defined on areal supports. Kyriakidis (2004) and Goovaerts (2008) show how the kriging system is adapted in the case of areal data supports and point prediction locations. *Area-to-point kriging* has theoretical advantages over the methods outlined previously in that it explicitly accounts for spatial structure in the variables.

3 | GRID-BASED POPULATION DATASETS

All areal data are subject to the modifiable areal unit problem (MAUP) whereby the results of analyses are a function of the size and shape of areal units (Openshaw, 1984; Openshaw & Taylor, 1979; Wong, 2009; see Weir-Smith, 2016 for a discussion in a South African context). However, with grids, the analyses are simplified as all units are of the same size and shape, and scale effects can be explored through simple aggregation of cells. In addition, a population grid “smooths” out spatial population discontinuities which are an artefact of the underlying arbitrary statutory geographies. Population grids are generated as standard outputs from censuses in many individual countries, including Estonia, Finland, the Netherlands and Sweden (see Batista e Silva et al., 2013; Gallego, 2010). There are several initiatives which have sought to develop gridded population or built up area datasets on a global basis. These include the Global Human Settlement Layer (GHSL; see Pesaresi et al., 2016) which is supported by the European Commission and provides gridded data on both built-up areas and population (<http://ghsl.jrc.ec.europa.eu/partners.php>). The GHSL population grid data are available at 250 m and 1 km spatial resolutions at a global scale for 1975, 1990, 2000 and 2015. The GHSL data for South Africa have been used as a backdrop for this research. The WorldPop project (see Tatem, 2017; Wardrop et al., 2018) is producing a wide array of gridded datasets, with a focus on Central and South America, Africa and Asia (<http://www.worldpop.org.uk/>).

Section 2 details some approaches which have been developed to reallocate data on irregular areal units to regular grids. Grids have been generated using a diverse array of additional information sources to inform reallocation of population counts from source zones (e.g., SALs) to target zones (e.g., 250 m by 250 m grid cells). The approaches employed differ depending on the quality of the input data. In cases where detailed household location data (with multiple attributes per household) and geographically rich census data (or other population surveys) are available, simple areal weighting and land use data may be sufficient to produce accurate population estimates for grid cells. In the present case, we have household locations but not detailed attributes per household. Remotely sensed imagery is used in many population gridding initiatives, while existing land use classifications are another key source of information for creating population grids. The WorldPop grids were generated using an array of approaches and data sources² – these include the random forest regression tree-based mapping approach detailed by Stevens et al. (2015), which incorporates information from multiple sources including, for example, remotely sensed imagery on night time lights, topography, land use and climate data. The GHSL grids were generated using remotely sensed imagery, national censuses, and also volunteered geographic information.

4 | STUDY AREA AND DATA

Census data at SAL level for 2001 and 2011 (see StatsSA, 2012) for Cape Town provide the basis for the analysis (see Mokhele et al., 2016 for more details). SALs are combinations of smaller zones called enumeration areas; there were 84,907 SALs in South Africa in 2011 with a mean population (using 2011 census counts) of 540 people. The example of unemployment is used in this paper. Recognising uncertainties in the data values due to small numbers at SAL level, a shrinkage approach was used. Shrinkage estimation is used to “borrow strength” from larger areas (in this case, local municipalities) to reduce the uncertainty associated with small area data (Noble et al., 2006). The end result of shrinkage is intended to move a SAL’s values towards a more reliable higher-level value which, in relation to the present example, might mean an adjustment towards either greater or lesser unemployment levels. Details of the shrinkage approach are provided by Smith et al. (2015). The allocation of counts (after application of shrinkage) from SALs to grids was undertaken using several data sources. These include Spot Building Count (SBC) data produced by ESKOM and the Council for Scientific and Industrial Research (CSIR; Ngidi et al., 2017) and Open Street Map (OSM) landuse data (OpenStreetMap contributors, 2017). The SBC data were developed using SPOT satellite imagery, and the dataset is intended to include all classifiable building structures within South Africa (Breytenbach, 2010). Each of the SBC points is linked to a potential population determined through overlay with sub-place census data – the SBC points are joined to sub-places using a GIS and the sub-place population allocated to the corresponding SBC points.

5 | GENERATION OF POPULATION SURFACES FOR CAPE TOWN

Area-to-point kriging is illustrated using the example of unemployment in Cape Town. Poisson kriging is a variant of kriging which is well suited to the analysis of population characteristics where rates may, in some cases, be computed from small numbers; this approach provides the basis of the present analysis. In this case study, census data for the years 2001 and 2011 were released for small areas which differ for each time period. To explore neighbourhood change, it is necessary to reallocate counts from the original source geographies to a set of common geographies; the data for 2001 and 2011 are thus reallocated from SALs to a 250 m grid. Figure 3 shows the percentage of the population who were unemployed by SAL for Cape Town in 2011. A 250 m grid was selected for two key reasons: it allows for analysis of spatial inequalities at a sufficiently fine spatial level and it matches the GHSL data introduced above and which have been used as a base layer in the project of which this work is part.

The population grids were generated using an input grid which indicates populated areas. This input grid was created using several stages:

1. Use the OSM mapping layers to produce a new layer which includes areas we can be confident are *not* residential, which we term the “remove” layer.
2. Remove all SBC data points that fall within polygons flagged as non-residential in the “remove” layer.
3. Remove all SBC data points that fall outside of study area of the City of Cape Town metropolitan municipality, and remove all SBC data points that do not lie within population census SALs (since within Cape Town there are some gaps in the 2011 SAL layer, indicating these areas are not populated).
4. Create empty 250 m grid.
5. Overlay refined SBC data points with 250 m grid and compute sum of SBC data population for SBC points within 250 m cells.
6. Keep only 250 m cells with a summed population >0.5 (rounded to a whole unit – one person).

This results in a layer of grid cells which we assume represents the coverage of population distribution across the City of Cape Town metropolitan municipality. The resulting grid cell centres are then used as targets with area-to-point kriging – in other words, unemployment rates are estimated for each of the grid cells.

5.1 | Area-to-point kriging

As a first stage, variograms of unemployment percentages were computed. Figure 4 shows the experimental areal variogram of unemployment % for SALs for 2011. The variogram exhibits two breaks of slope – the first at approximately 2.5 km, and the second at approximately 17.5 km; these correspond to city-wide (smaller figure) and region-wide (larger figure) spatial structures in the data. Figure 3 is the standard areal variogram; the next stage of the analysis was based on the variogram of the

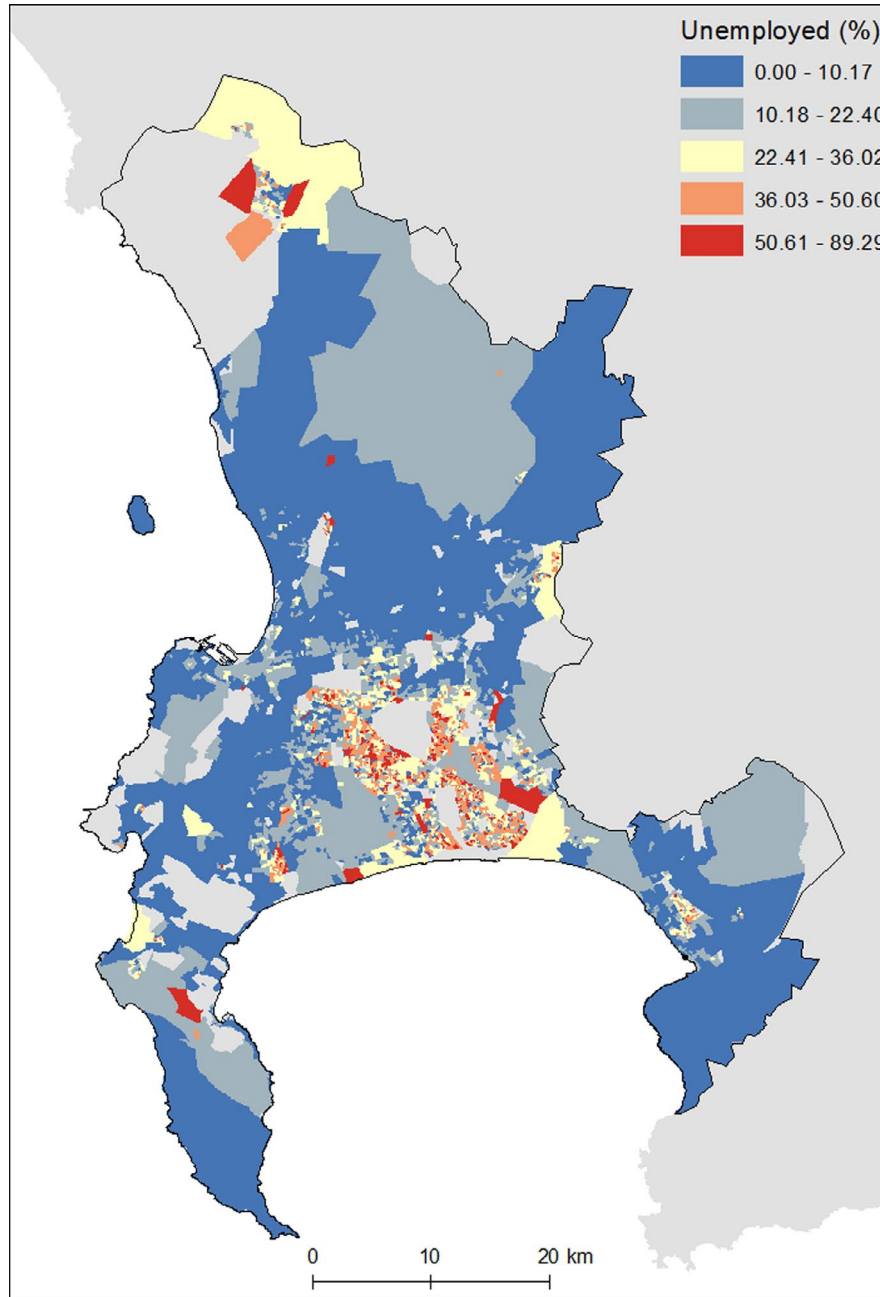


FIGURE 3 Unemployed persons (%) in SALs in 2011 in Cape Town

unknown risk. Goovaerts et al. (2005) characterise spatial variation in cancer mortality risk (a “rare” event – thus rates derived from these data and the total numbers of events are small). In cases like this, it is necessary to account for the reliability of observations – a function of population size. The variogram of the unknown risk is estimated following Goovaerts et al. (2005). Area to point kriging depends on the point variogram, which was derived from the areal variogram of the unknown risk using deconvolution, as introduced above. Figure 5 shows the model fitted to this areal variogram and also the model derived using deconvolution. As expected, the deconvolved model has a larger sill than the areal model and this reflects the objective of deconvolution – to account for the variation lost by aggregation to zones such as, in this example, SALs.

The deconvolved model is used next to inform estimates with area-to-point kriging. Ordinary kriging with Poisson population (total population; here total employed and unemployed people) adjustment was applied with a population denominator of 1 (i.e., the analysis is based on proportions). The discretisation geography was 250 m cells. The destination geography (locations where estimates are required) was 250 m cells (as for the discretisation geography). The search neighbourhood was a quadrant using a minimum of 1 and a maximum of 16 observations with a search radius of 28.9 km (these figures were derived




FIGURE 4 Experimental variogram of unemployment % for SALs in Cape Town in 2011

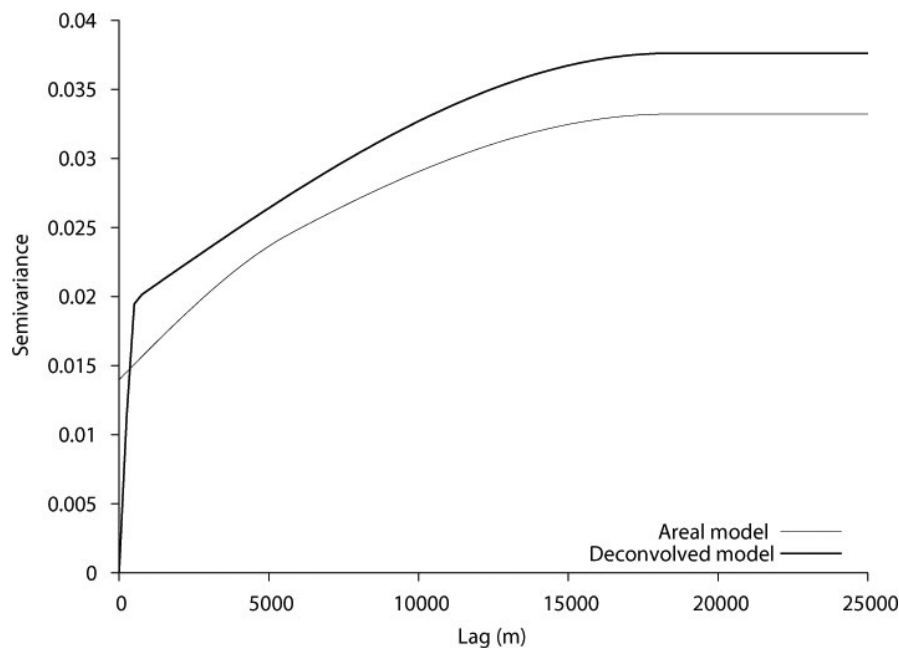


FIGURE 5 Model fitted to Poisson variogram and deconvolved model for unemployment rates in Cape Town in 2011

experimentally and account for the average spacing between SAL centres). The unemployment rate estimates for 2011 are shown in Figure 6(a) while the percentage point changes in unemployment (2011 estimates minus 2001 estimates) are shown in Figure 6(b). A key benefit in the gridded estimates is that there are “holes” where there are no people. This allows for more accurate depiction of, for example, spatial inequalities since, with standard zonal data, using shared boundaries as a measure of likelihood of interaction may be flawed if the area covered by the shared boundary is unpopulated. Where this process is completed for several time points with incompatible zonal systems, it becomes possible to explore local changes. It is worth noting that in some areas, and especially in the north of the region, standard zones would suggest there is considerable homogeneity in population characteristics whereas, with 250 m grid cells, the mixed characteristics and sparsity of the population in these areas is clearly apparent.

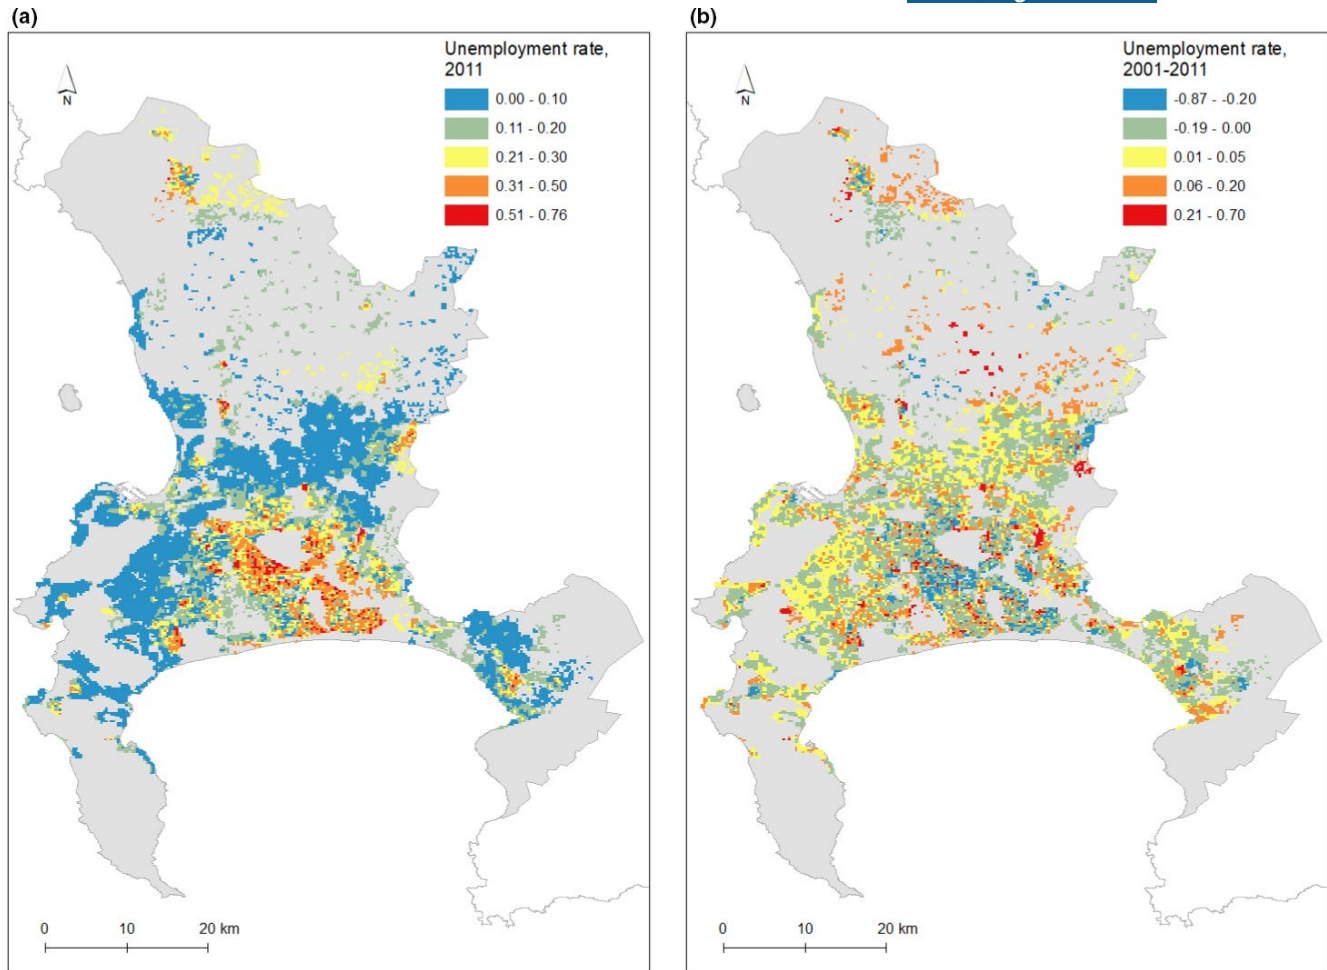


FIGURE 6 (a) Unemployment rate in Cape Town in 2011. (b) Percentage point change in unemployment between 2001 and 2011 (positive differences are increases; negative differences are decreases)

5.2 | Quality of the gridded predictions

There is no direct means to assess prediction accuracy in this case as there is no existing gridded dataset to which the derived grids can be compared. However, previous work in other national contexts may be informative. Lloyd et al. (2017) used an areal weighting approach with a simple smoothing procedure and this was assessed using gridded population data for Northern Ireland. In this case, the largest population errors (predicted values minus observed values) were found in areas where high-density populations (for example, in tower blocks) were wrongly “spread” across larger areas. The main focus in the present paper is on rates rather than counts, and the strong spatial structure of deprivation in most areas of Cape Town suggests that splitting source zones into (smaller) gridded cells, as is done here, will lead to spatially accurate distributions since kriging disaggregates to smaller areas while accounting for differences between neighbouring areas. This is further demonstrated later using a case study in an area of Cape Town called Dunoon. Figure 1 showed an example of the gridded unemployment rates based on the kriging approach superimposed on SAL boundaries – gridded unemployment rates within a given SAL based on a simple apportionment approach would be identical. The differences are largest in areas with spatially contrasting unemployment values (i.e., neighbouring areas are very different) and smallest in areas with similar unemployment levels. This is conceptually sensible as it reflects the “transition” between areas with high and low unemployment, but retains the distinction between areas at their borders.

6 | ANALYSIS OF SPATIAL INEQUALITIES

A key reason for developing gridded population variables is the need to chart how spatial inequalities have changed across small areas of Cape Town. One approach to measuring spatial inequalities is to compute a measure of spatial autocorrelation. One of the most widely applied measures of spatial autocorrelation is the *I* coefficient (Moran, 1950). The *I* coefficient measures covariation in a single variable measured at multiple locations. An example would be deprivation levels in census areas, and the concern is with assessing how far neighbouring deprivation values tend to be similar. First, we define a neighbourhood – with a regular grid, we could simply say that all zones which share edges or corners with other zones are neighbours of that zone; this is termed queen contiguity. Moran's *I* is given by:

$$I = \frac{n \sum_{i=1}^n \sum_{j=1}^n w_{ij} (y_i - \bar{y})(y_j - \bar{y})}{\left(\sum_{i=1}^n (y_i - \bar{y})^2 \right) \left(\sum_{i=1}^n \sum_{j=1}^n w_{ij} \right)} \quad (2)$$

The right hand part of the numerator, $w_{ij}(y_i - \bar{y})(y_j - \bar{y})$, comprises the weights for paired data locations *i* and *j* multiplied by the covariance between y_i and y_j – the mean is subtracted from each value and the products are multiplied. The sum of these covariances for all paired locations is multiplied by *n* – the number of observations. The output is divided by the sum of the squared differences between all of the data values and their mean average multiplied by the sum of all of the weights. In the case of queen contiguity, the weights for all individual neighbours of a zone would be one. In many applications, row standardisation is used and the weights are divided by the number of neighbours. As an example, if there are five neighbours, then the weights all become $1/5 = 0.2$. Moran's *I* for unemployment for a 250 m grid for 2001 was 0.816 (pseudo $p = .0001$), while the equivalent figure for 2011 was 0.702 (pseudo $p = .0001$). This suggests that neighbouring areas have, on average, become *less* similar by unemployment. This could suggest that spatial inequalities have increased – in other words, neighbouring areas are less likely to have similar unemployment levels. This can be interrogated further through a local approach. Various local measures of spatial autocorrelation have been developed. One of the most widely used is a local variant of Moran's *I* presented by Anselin (1995). It is given by:

$$I_i = (z_i/s^2) \sum_{j=1}^n w_{ij} z_j, j \neq i \quad (3)$$

where z_i are differences of variable *y* from its global mean ($y_i - \bar{y}$). In cases where zones are used (as opposed to points), the weights, w_{ij} , are often set to 1 for immediate neighbours of a zone and 0 for all other zones (queen contiguity); s^2 is the sample variance. Note that local *I* values sum up to global Moran's *I*. If a zone has a large percentage of group *m* (e.g., unemployed people) and it has several neighbouring zones with very large percentages of group *m* (values which are larger than average levels), then the value of *I* will be large and positive. If a zone has a very large percentage of a particular group (larger than the average) but its neighbours have very small percentages (smaller than the average), then the value of *I* will be large and negative. It is worth noting that the results are a function of the MAUP, and the results obtained using Sub-Place data rather than SALs as source zones would be different; that is, the larger the source zone, the less variation there would be between constituent grid squares. It is also worth noting that queen contiguity is a very narrow definition of what constitutes neighbouring areas, especially when using a fine grained 250 m grid square geography, but is used here as a simple definition of a local neighbourhood.

Anselin (1995) describes an approach to testing for significant local autocorrelation based on random relocation of the data values, the objective being to assess if the observed configuration of values is significant. Significant clusters of values can be computed using local *I*. The end result is a set of five categories: high-high (large values [proportions] surrounded by large values), low-low (small values surrounded by small values), high-low (large values surrounded by small values), low-high (small values surrounded by large values), and not significant. The GeoDa software offers the capacity to test the significance of local *I* using randomisation³ (this can be used to derive pseudo *p* values) and to map significant clusters. Clusters are identified using the Moran scatterplot (Anselin, 1995).

Figure 7 shows significant clusters in the unemployment rate in Cape Town in (a) 2001 and (b) 2011. There are very distinct spatially continuous areas comprising low-low clusters (areas with low levels of unemployment surrounded by areas with similar characteristics), and large areas (principally in a zone to the east of the city centre known as the Cape Flats, which includes a number of townships, such as Khayelitsha), comprising high-high clusters. Within the latter area, there are a number of neighbourhoods with lower levels of unemployment – these are represented as low-high clusters. The broad patterns are consistent across the two time points, but there are obvious transitions from, in particular, high-high to low-high clusters.

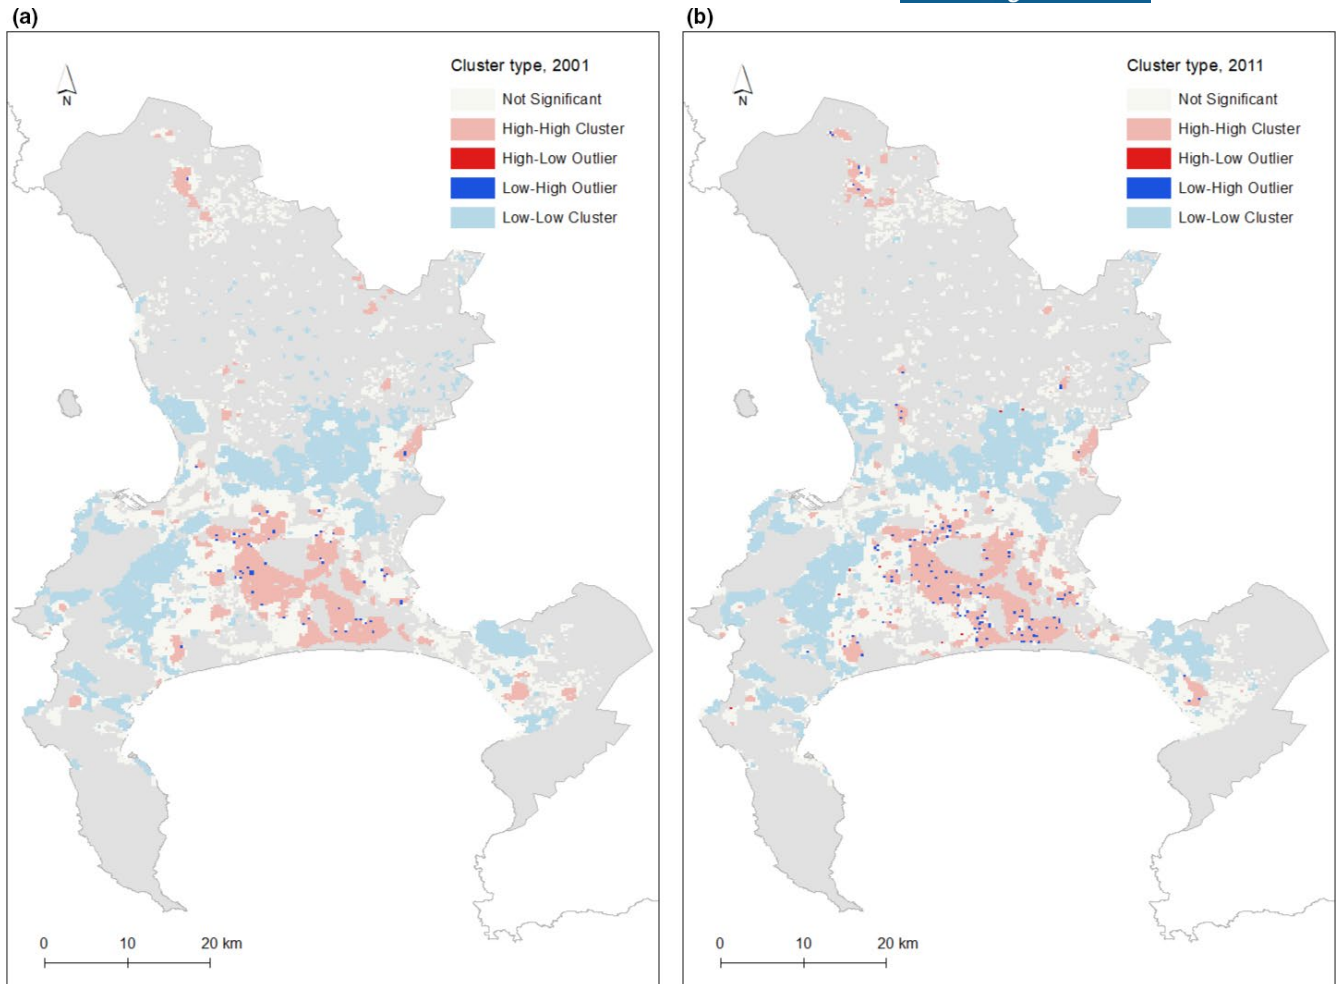


FIGURE 7 Significant clusters in unemployment rate in Cape Town in (a) 2001. (b) 2011

These correspond to areas which have reduced levels of unemployment relative to neighbouring areas and in these cases, spatial inequalities have increased.

6.1 | Example: Dunoon

Here, the region around Dunoon (located some 10 km north of the centre of Cape Town) provides a specific focus to assess changing unemployment patterns. Context to Dunoon is provided by McGaffin et al. (2015); in brief, the establishment of Dunoon in 1995 took place under the Less Formal Township Establishment Act (LEFTEA) legislation. Dunoon makes a useful case study as it is a deprived neighbourhood surrounded by relatively affluent areas; these include Parklands, Table View and Killarney Gardens (October & Freeman, 2017). There is much informal housing (according to the 2011 census, 41% of households lived in formal dwellings), population density in the area is high, and service provision (including sewerage and bin collection) is generally poor. Most occupants rent their homes. One feature of the area are “backyarders” – people who pay rent to reside in the backyards of other people’s homes (October & Freeman, 2017) and in some cases these are rented from people who are themselves renters. This has implications for service delivery, as the municipality requires that property owners must be present in order to make complaints about service provision (October & Freeman, 2017). Dunoon was also subject to an outbreak of violence aimed at foreigners during a series of xenophobic attacks across South Africa, commencing in 2008 (Cooper, 2009; October & Freeman, 2017).

Figure 8 shows clusters for 2001 and 2011 unemployment rates centred on the Dunoon region with satellite imagery used as a backdrop. The satellite imagery highlights that grid cells only exist for populated areas – with fields and industrial buildings shown outside of the area covered by the grid cells. In the map for 2001, the central area includes grid cells which are categorised as high-high clusters, neighboured by grid cells categorised as not significant. In the equivalent map for 2011, some cells




FIGURE 8 Significant clusters in unemployment rate in the Dunoon area in (a) 2001. (b) 2011. SAL boundaries are shown in red
Source: The backdrop is the Google satellite layer. Map data ©2018 Google.

have transitioned from not significant to high-high, while a small number have transitioned from not significant, or high-high to low-high. There are, therefore, changing patterns of spatial inequality.

The focus of McGaffin et al. (2015) was on the vertical consolidation of Dunoon, a relatively rare case in South Africa of the demolition of state-built structures by property owners and their replacement by double-storey rental accommodation. This development of new rental accommodation may reflect growing spatial inequalities between those in newly built property with provision for rental and those in the original state-developed housing. This cannot be directly assessed in this case study given its focus on unemployment specifically, but it presents an interesting case for future work on changing patterns of deprivation generally. Examination of grids of total population values show that most of Dunoon is associated with increasing population densities over the time period. It seems likely that this increase is linked to growth, in at least some areas, in informal housing. The transitions from high-high to low-high clusters are indicative either of relative improvement in the grid cells concerned (as suggested above, possibly connected to replacement of state built housing with new rental accommodation), or a relative worsening of conditions in the neighbouring cells. There is considerable scope to unpick these patterns at a local level, but it is outside the remit of this study.

7 | DISCUSSION AND CONCLUSIONS

The grid generation approach is being applied to a host of population variables for each of the South African Censuses of 2001 and 2011, for the whole of South Africa. The key focus of this work is on the construction of deprivation measures based on income poverty, employment, education, and living environment for each grid cell. This will build on previous work which has sought to measure multiple domains of deprivation across South Africa (Noble & Wright, 2013; Noble et al., 2013). The end result of this programme of work will be a major resource for charting changes in deprivation and spatial inequalities across the country. By making links to qualitative work undertaken in Cape Town it will be possible to infer the possible impacts “on the ground” of these changing inequalities. The potential social and policy impact of this resource is considerable offering, as it will, a comprehensive overview of where inequalities are largest, where they have changed most, and in what ways. An understanding of the population and economic trajectories and of the implications of spatial inequalities between areas (see Sinclair-Smith & Turok, 2012 for a related study of Cape Town) is vital in developing schemes to reduce these inequalities.

Only one measure of spatial inequality is used in this paper – the Moran's *I* spatial autocorrelation coefficient. Spatial inequality has links to segregation and the use of a suite of local measures of segregation to capture different facets of spatial inequality could be assessed. Previous work suggests that the exposure dimension may relate more strongly to people's lived experiences of inequality than other dimensions of segregation (see McLennan et al., 2016). As well as the dimension of segregation, definitions of neighbourhood are important. In this paper, spatial clustering is measured using adjacent grid cells which, it may be argued, is a relatively crude definition of neighbourhood. Other possibilities are distance decay functions whereby the likelihood of interactions between people in areas are assumed to reduce as distance between areas increases. More sophisticated ways of measuring potential interactions (and thus exposure to people with different socio-economic characteristics) include the use of cost surface analysis, as applied by Lloyd (2015) in an analysis of residential segregation by religion in Belfast. Future work will assess some such schemas in the context of South Africa.

The availability of data on a consistent geographical basis allows for the assessment of the importance of persistent or rapidly changing spatial inequalities – a much under-explored area of research. The findings from this quantitative analysis have been linked to a qualitative phase of the work using the South African Social Attitudes Survey and geographically targeted focus groups. This work has sought to assess the views of respondents on their own lived experience of inequalities, building on the work of McLennan et al. (2014). This element of the research is exploring the associations between measured levels of inequality and attitudes to inequalities and provisional results show that the relationships between the two are complex.

The value of gridded population data is recognised by StatsSA. Verhoef (2019) details options for population data output geographies and this includes discussion of the Basic Spatial Unit (BSU) frame comprising a 100 m grid which will provide the basis for harmonising data from different sources, to allow for exploration of changes across time, and potentially, also comparisons between countries. There is, therefore, the potential to link the grids detailed in this paper with those which will be produced by StatsSA in the future – albeit with a need to create grids with the same spatial resolution. This would provide a powerful means to assess how inequalities have changed over small areas and to consider the determinants of change in these areas by linking to the rich array of variables available in the census.

The grids described in this paper are accessible via the project website,⁴ both as raw data files and via dynamic mapping interface. The website also provides access to deprivation measures for all of South Africa for census wards which are consistent for 2001 and 2011, thus allowing direct measurement of change between these two time points. In addition, supporting

documentation on the creation of the grids and methods for their analysis are available. The resource includes a range of measures of spatial inequalities, including measures of exposure (see McLennan et al., 2016), for both grids and standard geographical zones. This is supported by guidance on how to manipulate and analyse the data using the open source GIS package QGIS (QGIS Development Team, 2020).

In the present analysis, the sole focus is on data which are specific to South Africa. However, the availability of population grids globally means that the principles can be applied anywhere where areal data exist (e.g., from a census) and the desire is to reallocate the counts to a regular grid. The paper contributes to wider body of work which seeks to utilise diverse data sources to more accurately map population distributions – a key example of this is the WorldPop project (see, for example, Tatem, 2017). The present paper is a novel contribution to the literature on spatial inequalities. The use of population grids could become the cornerstone of such analyses internationally, and particularly where there is a desire to assess the nature of spatial inequalities between countries or regions.

ACKNOWLEDGEMENTS

InequalitiesSA was funded by the Economic and Social Research Council (ESRC) and the National Research Fund (NRF) under the ESRC-NRF Newton Call for Collaborative Research Urban Transformations in South Africa call (project ES/N014022/1), supported by the Newton Fund. This support is acknowledged gratefully. The project builds upon an earlier ESRC Pathfinder project which was undertaken by members of the core research team: “The relationship between spatial inequality and attitudes to inequality in post-apartheid South Africa” (project ES/I034889/1). Map data copyrighted OpenStreetMap contributors and available from <https://www.openstreetmap.org>. We are grateful to Statistics South Africa for provision of Census data for 2001 and 2011. We thank the Council for Scientific and Industrial Research (CSIR) for provision of the SBC data. The authors thank the reviewers for their helpful comments which have helped to refine the paper.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available at <https://reshare.ukdataservice.ac.uk/854944/>

ORCID

Christopher D. Lloyd  <https://orcid.org/0000-0003-0575-7793>

ENDNOTES

¹ This compares to South Africa-wide figures of 0.72 in 2006 to 0.67 in 2015 for household per capita income (StatsSA, 2019).

² See <http://www.worldpop.org.uk/data/methods/> for an overview.

³ see <https://geodacenter.github.io/glossary.html>

⁴ <http://qub.ac.uk/research-centres/GIS/Research/InequalitiesSA/>

REFERENCES

- Anselin, L. (1995) Local indicators of spatial association — LISA. *Geographical Analysis*, 27, 93–115. Available from: <https://doi.org/10.1111/j.1538-4632.1995.tb00338.x>
- Atkinson, P.M. & Curran, P.J. (1995) Defining an optimal size of support for remote sensing investigations. *IEEE Transactions on Geoscience and Remote Sensing*, 33, 768–776. Available from: <https://doi.org/10.1109/36.387592>
- Atkinson, P.M. & Tate, N.J. (2000) Spatial scale problems and geostatistical solutions: A review. *Professional Geographer*, 52, 607–623. Available from: <https://doi.org/10.1111/0033-0124.00250>
- Batista e Silva, F., Gallego, J. & Lavalle, C. (2013) A high-resolution population grid map for Europe. *Journal of Maps*, 9, 16–28. Available from: <https://doi.org/10.1080/17445647.2013.764830>
- Breytenbach, A. (2010) A national inventory using earth observation and GIS. *Science Scope*, 5, 66–68. Available from: <https://doi.org/10.10520/EJC29689>
- Cooper, A. (2009) “Let us eat airtime”: Youth identity and ‘xenophobic’ violence in a low-income neighbourhood in Cape Town. Social Surveys Unit, Centre for Social Science Research. CSSR Working Paper No. 263. Centre for Social Science Research, University of Cape Town (UCT), South Africa (CSSR). Available at: <https://www.eldis.org/document/A45549>
- Gallego, J. (2010) A population density grid of the European Union. *Population and Environment*, 31, 460–473. Available from: <https://doi.org/10.1007/s11111-010-0108-y>
- Goovaerts, P. (2008) Kriging and semivariogram deconvolution in the presence of irregular geographical units. *Mathematical Geosciences*, 40, 101–128. Available from: <https://doi.org/10.1007/s11004-007-9129-1>
- Goovaerts, P., Jacquez, G.M. & Greiling, D. (2005) Exploring scale-dependent correlations between cancer mortality rates using factorial kriging and population-weighted semivariograms. *Geographical Analysis*, 37, 152–182. Available from: <https://doi.org/10.1111/j.1538-4632.2005.00634.x>

- Keeley, B. (2015) *Income inequality: The gap between rich and poor. OECD insights*. Paris, France: OECD Publishing.
- Kyriakidis, P.C. (2004) A geostatistical framework for area-to-point spatial interpolation. *Geographical Analysis*, 36, 259–289. Available from: <https://doi.org/10.1111/j.1538-4632.2004.tb01135.x>
- Lloyd, C.D. (2014) *Exploring spatial scale in geography*. Chichester, UK: John Wiley and Sons.
- Lloyd, C.D. (2015) Local cost surface models of distance decay for the analysis of gridded population data. *Journal of the Royal Statistical Society: Series A*, 178(1), 125–146. Available from: <https://doi.org/10.1111/rssa.12047>
- Lloyd, C.D., Catney, G., Williamson, P. & Bearman, N. (2017) Exploring the utility of grids for analysing long term population change. *Computers, Environment and Urban Systems*, 66, 1–12. Available from: <https://doi.org/10.1016/j.compenvurbsys.2017.07.003>
- Martin, D. (1989) Mapping population data from zone centroid locations. *Transactions of the Institute of British Geographers*, 14, 90–97. Available from: <https://doi.org/10.2307/622344>
- Martin, D. (1996) An assessment of surface and zonal models of population. *International Journal of Geographical Information Systems*, 10, 973–989. Available from: <https://doi.org/10.1080/02693799608902120>
- McGaffin, R., Cirolia, L.R. & Massyn, M. (2015) Overcoming the challenge of vertical consolidation in South Africa's low-income settlements: A case study of Du Noon. *Urban Forum*, 26, 59–75. Available from: <https://doi.org/10.1007/s12132-014-9232-6>
- McLennan, D., Noble, M. & Wright, G. (2016) Developing a spatial measure of exposure to socio-economic inequality in South Africa. *South African Geographical Journal*, 98, 254–274. Available from: <https://doi.org/10.1080/03736245.2015.1028980>
- McLennan, D., Roberts, B., Noble, M., Masilela, T. & Magidimish, H. (2014) *Exploring the relationship between spatial inequality and attitudes to inequality in South Africa*. Oxford, UK: Centre for the Analysis of South African Social Policy (CASASP), University of Oxford. Available at: <https://gtr.ukri.org/projects?ref=ES%2FI034889%2F1#tabOverview>
- Mokhele, T., Mutanga, O. & Ahmed, F. (2016) Development of census output areas with AZTool in South Africa. *South African Journal of Science*, 112(7–8), 1–7. Available from: <https://doi.org/10.17159/sajs.2016/20150010>
- Moran, P.A.P. (1950) Notes on continuous stochastic phenomena. *Biometrika*, 37, 17–23. Available at: <https://www.jstor.org/stable/2332142>
- Ngidi, M., Mans, G., McKelly, D. & Sogoni, Z. (2017) Using a hybrid methodology of dasyametric mapping and data interpolation techniques to undertake population data (dis)aggregation in South Africa. *South African Journal of Geomatics*, 6(2), 232–244. Available from: <https://doi.org/10.4314/sajg.v6i2.8>
- Noble, M., Wright, G., Smith, G. & Dibben, C. (2006) Measuring multiple deprivation at the small-area level. *Environment and Planning A: Economy and Space*, 38, 169–185. Available from: <https://doi.org/10.1068/a37168>
- Noble, M. & Wright, G. (2013) Using indicators of multiple deprivation to demonstrate the spatial legacy of Apartheid in South Africa. *Social Indicators Research*, 112, 187–201. Available from: <https://doi.org/10.1007/s11205-012-0047-3>
- Noble, M., Zembe, W., Wright, G. & Avenell, D. (2013) *Multiple deprivation and income poverty at small area level in South Africa in 2011*. Cape Town, South Africa: SASPRI. Available at: http://www.saspri.org/wpcontent/uploads/Docs/SAIMD_2011_ward_level_National_Report_for_Web.pdf
- October, L. & Freeman, L. (2017) *Dunoon case study. Xenophobia and outsider exclusion: addressing frail cohesion in South Africa's diverse communities*. Cape Town, South Africa: University of Cape Town. Available at: https://freedomhouse.org/sites/default/files/South_Africa_Community_Social_Cohesion_Profiles_Synthesis%20Report.pdf
- Openshaw, S. (1984) *The modifiable areal unit problem. Concepts and techniques in modern geography* 38. Norwich, UK: GeoBooks.
- Openshaw, S. & Taylor, P.J. (1979) A million or so correlation coefficients: Three experiments on the modifiable areal unit problem. In: Wrigley, N. (Ed.) *Statistical applications in the spatial sciences*. London, UK: Pion, pp. 127–144.
- OpenStreetMap Contributors. (2017) *Planet dump* [Data file from 2017]. Available at: <https://planet.openstreetmap.org>
- Pesaresi, M., Ehrlich, D., Ferri, S., Florczyk, A.J., Freire, S., Halkia, S. et al. (2016) *Operating procedure for the production of the global human settlement layer from landsat data of the epochs 1975, 1990, 2000, and 2014*. Publications Office of the European Union, EUR 27741 EN, 2016. <https://doi.org/10.2788/253582>
- QGIS Development Team. (2020) QGIS geographic information system. Open Source Geospatial Foundation Project. Available at: <http://qgis.osgeo.org>
- Schabenberger, O. & Gotway, C.A. (2005) *Statistical methods for spatial data analysis*. Boca Raton, FL: Chapman and Hall/CRC Press.
- Sinclair-Smith, K. & Turok, I. (2012) The changing spatial economy of cities: An exploratory analysis of Cape Town. *Development Southern Africa*, 29(3), 391–417. Available from: <https://doi.org/10.1080/0376835X.2012.706037>
- Skøien, J.O., Blöschl, G., Laaha, G., Pebesma, E., Parajka, J. & Viglione, A. (2014) rtop: An R package for interpolation of data with a variable spatial support, with an example from river networks. *Computers and Geosciences*, 67, 180–190. Available from: <https://doi.org/10.1016/j.cageo.2014.02.009>
- Smith, T., Noble, M., Noble, S., Wright, G., McLennan, D. & Plunkett, E. (2015) *The English indices of deprivation 2015 technical report*. London, UK: Department for Communities and Local Government. Available at: <https://www.gov.uk/government/publications/english-indices-of-deprivation-2015-technical-report>
- StatsSA. (2012) *Census 2011 statistical release*. Pretoria, South Africa: Statistics South Africa. Available at: <https://www.statssa.gov.za/publications/P03014/P030142011.pdf>
- StatsSA. (2019) *Inequality trends in South Africa: a multidimensional diagnostic of inequality*. Report No. 03-10-19. Pretoria, South Africa: Statistics South Africa. Available at: <http://www.StatsSA.gov.za/publications/Report-03-10-19/Report-03-10-192017.pdf>
- Stevens, F.R., Gaughan, A.E., Linard, C. & Tatem, A.J. (2015) Disaggregating census data for population mapping using random forests with remotely-sensed and ancillary data. *PLoS One*, 10(2), e0107042. Available from: <https://doi.org/10.1371/journal.pone.0107042>
- Tatem, A.J. (2017) WorldPop, open data for spatial demography. *Scientific Data*, 4, 170004. Available from: <https://doi.org/10.1038/sdata.2017.4>

- Verhoef, H. (2019) Census geography: What constitutes the ideal output area for South Africa? *PositionIT*, 30 July. Available at: <https://www.ee.co.za/article/census-geography-what-constitutes-the-ideal-output-area-for-south-africa.html>
- Wardrop, N.A., Jochem, W.C., Bird, T.J., Chamberlain, H.R., Clarke, D., Kerr, D. et al. (2018) Spatially disaggregated population estimates in the absence of national population and housing census data. *Proceedings of the National Academy of Sciences of the United States of America*, 115(14), 3529–3537. Available from: <https://doi.org/10.1073/pnas.1715305115>
- Webster, R. & Oliver, M.A. (2007) *Geostatistics for environmental scientists*, 2nd edition. Chichester, UK: John Wiley & Sons.
- Weir-Smith, G. (2016) Changing boundaries: Overcoming modifiable areal unit problems related to unemployment data in South Africa. *South African Journal of Science*, 112(3/4), 1–8. Available from: <https://doi.org/10.17159/sajs.2016/20150115>
- Western Cape Government Provincial Treasury. (2012) *Regional development profile city of Cape Town*. Working paper. Cape Town, South Africa: Western Cape Provincial Treasury. Available at: [https://www.westerncape.gov.za/text/2011/12/dc0_city_of_cape_town_15_dec_2011_\(mb\).pdf](https://www.westerncape.gov.za/text/2011/12/dc0_city_of_cape_town_15_dec_2011_(mb).pdf)
- Wong, D. (2009) The modifiable areal unit problem (MAUP). In: Fotheringham, A.S. & Rogerson, P.A. (Eds.) *The SAGE handbook of spatial analysis*. London, UK: SAGE Publications, pp. 105–123.

How to cite this article: Lloyd, C.D. Bhatti, S. McLennan, D. Noble, M. & Mans, G. (2021) Neighbourhood change and spatial inequalities in Cape Town. *The Geographical Journal*, 00, 1–16. <https://doi.org/10.1111/geoj.12400>