On-Line Appendices

Appendix 2: Data Methodology

This Appendix has three components. The first discusses and describes the sources for all data used in this paper. The second deals with measures on cover/footprint and volume we use to analysis. The third gives the algorithm used to extract unchanged buildings, redeveloped buildings and infill from the overlay of 2004 and 2015 depiction of building polygons.

Data Sources

Building Data

We use two cross sections of data that delineate every building footprint in the city of Nairobi. The first is based on tracings of buildings from aerial photo images for 2003 which we received from the Nairobi City Council. Although no explicit metadata was provided, as far as we can tell this data was created by the Japan International Cooperation Agency (JICA) and the Government of the Republic of Kenya under the Japanese Government Technical Cooperation Program, and based on aerial images taken in February 2003 at a scale of 1:15,000. We base this off documentation from the Center for Sustainable Urban Development (CSUD) at Columbia University, who use a highly detailed building density and landuse map from the JICA (Williams et.al. 2014). Further we do our own data quality check by comparing the digital tracings to very high resolution imagery from Google Earth (2002), (2003), and (2004). By examining areas that changed from 2002-2003 and from 2003-2004 we confirm that our data of building outlines matches those that exist in 2003, but did not exist in 2002, and does not include those that were yet to be built in 2003 and appeared in 2004. The second cross section comes from January 2015, when imagery at (10-20cm resolution) was recorded and digitized into building footprints by a Nairobi based company Ramani Geosystems.

The footprint data describe only the area on the ground that each building occupies while we are interested in the complete volume of each building. To address this need we supplement the 2-dimensional building data with 2015 building height data derived from LiDAR (0.3-1m resolution) which was again produced by Ramani Geosystems. Without direct measurements of heights in 2003, we interpolate them by assigning to each building in a grid square in a sector (slum or formal) the average height of unchanged buildings in the same sector over queen neighbouring grid squares.

Slum and land use maps

We focus on a definition of slums provided IPE Global under the Kenya Informal Settlements program (KISIP). IPE mapping of informal settlements was done using satellite imagery and topographic maps. Their approach was to identify slums as “unplanned settlements” which have some aspects of low house quality, poor infrastructure, or insecure tenure. To incorporate this definition of slums into our database we created shapefiles by manually digitizing KISIP documentation which contained detailed maps of all identified informal settlements in Nairobi (IPE Global Private Limited and Silverwind Consultants, 2013). There remains an issue of tight delineation of slum areas, where boundaries are drawn to outline the slum areas leaving a lot of empty land residual in the formal sector which we define as the complement to slums. To offset this, we adjust the IPE slum boundaries by first classifying buildings as slum if their centre lies within the original slum boundary, and then assigning each 3m x 3m pixel of non-built land to slum if the nearest building is classified as slum, and formal otherwise.

A secondary set of maps that we use comes from the Center for Sustainable Urban Development (CSUD) at Columbia University. The CSUD maps landuse in 2003, including slums, based on a more detailed, copyrighted, land-use map created by the JICA and the Government of Kenya under the Japanese Government Technical Cooperation Program which was published and printed by the survey of Kenya 1000 in March 2005 (Williams, et al. 2014). In principle, polygons are categorized as slums if they seemed to contain small mostly temporary buildings that are randomly distributed in high density clusters. We use this set of slums to offer a descriptive comparison of how slums have changed on the extensive margin, but for our analysis we defer to a single definition based on IPE due to discrepancies in the definition of slum across the data sources. We also make use of the CSUD landuse map to identify areas that we remove from our formal classification. The areas that we chose to remove are listed in appendix table 3.3 and are areas in permanent public use.

Household Survey

In order to get estimates on slum and formal household rents we use a cross section of georeferenced household level data from the 2012 'Kenya: State of the Cities' survey by the National Opinion Research Center (NORC) (Zinnes et.al. 2012). This is the first survey to record *household* rent (with detailed house and some neighbourhood characteristics) for a sample that is stratified between slum and formal areas (based on the 2009 Census) covering Nairobi. Also included in this survey were geo-coordinates taken at the time of survey, however we found these to be imprecise when compared to the location of the enumeration area (EA) that the household was recorded to reside in. We correct household coordinates if they fall outside of their EA by replacing them with the EA's centroid coordinates.

Vacant land price listings

We also require data on land values in order to calibrate the model, for this we rely on property values that have been scraped from property24.co.ke over the period September 2014 to November 2015. This data source provides us with vacant land listings recording information on asking price and plot area and location, all of which are provided for in over 80% of the listings. The locations are descriptive and so we entered geo-coordinates by manually searching the addresses and location descriptions. These listings are only found in the formal sector.

SRTM elevation

Elevation and ruggedness measures used in regression tables are calculated from the Shuttle Radar Topography Mission (SRTM), a grid of 1 arc-second wide cells (or roughly 30 metres in Nairobi) published by the USGS (2005). Elevation is simply the mean of these cells in each of our 150x150m gridcells, while we measured ruggedness as the standard deviation in elevation within each 150x150 metre gridcell.

SPOT Imagery

We also use high resolution SPOT5 and SPOT6 images of Nairobi for 2004 and 2015 respectively. The raw imagery was created by Airbus Defence and Space and we used it as reference to manually trace roads and define their widths in order to come up with estimates of the extent of road coverage in both the early and late time periods. Alternative sources, like Open Streetmap, were unsuitable as they did not allow us to make the comparison across time.

Measures of cover and volume

Our unit of analysis is 150x150m grid squares. For calculating cover within the grid square in a usage, each of these is broken into 50 3x3m cells and use type classified by what is at the centroid of the 3m square in each period. There are three uses: vacant land, slum area and formal. For each 150x150 square we sum across the 50 cells to get total use of each type. Most 150x150 squares are either all slum or all formal sector. However there are about 12% which are mixed grid squares, for which we record the cover or volume of slum and formal separately.

Having summed the total area of use of each type in 3x3 squares in each 150x150 meter square, these are averaged for 150x150m squares whose centroid falls in a narrow distance ring. That sum is then divided by the total number of 150x150 grid squares in that distance band. For volume for 2015, for each 3x3m square which is formal sector, we have the height of the building at the centroid of that square. Volume for that 3x3 square is 9 times the height in meters of the building from LiDAR data. We then sum across the grid squares occupied with formal usage for 150x150m grid squares in each distance ring and then average by the total number of 150x150 m grid squares in the ring. For 2004 we have no height data. To infer 2004 heights, we use what we think is an upper bound on height: the height of unchanged buildings, where we presume demolished buildings between 2004 and 2015 are likely to be of lower height than those which survive. To assign a height to a 3x3m square in 2004 in formal sector usage, we take the average height in 2015 of all buildings that were there in 2004 for all 3x3m formal sector unchanged buildings in the own 150x150m grids square and its 8 queen neighbours. Height is the height assigned to each 3x3m square in usage in a distance ring from the centre averaged over all such cells, to effectively get a coverage weighted average of individual building heights.

How do we measure change between 2004 and 2015? For demolition, at the 3x3m level the square is defined as demolition if its centroid is covered by a 2004 building which has been replaced by open space. Demolished coverage is lost 2004 cover; demolished volume is assessed as before using the average height of unchanged buildings in the neighbourhood. Infill is new buildings which do now overlap with any 2004 buildings; a 3x3m square is infill if its centroid is covered by such a building on 2015 where there was no building in 2004. Infill cover and volume are assessed from 2015 data. Net redevelopment in coverage takes coverage in the new 2015 buildings and subtracts the coverage of old 2004 buildings. So for each 150m150m meter square we have for redeveloped buildings, we have total coverage in 2004 measured at the 3x3m level (centroid covered by the old 2004 building(s)) and we have total coverage in 2015 measured at the 3x3m squares (centroid covered by the new replacement 2015 building(s)). Net redevelopment at the 150x150sqaure is the difference. In general, the same buildings are drawn in 2015 to have modestly more coverage than in 2004 so coverage change is likely to be an upper bound. Net volume change again assigns heights in 2004 to the 3x3m coverage based on neighbourhood averages for unchanged buildings and uses 2015 height information on the new buildings.

Overlaying Buildings

We match buildings across time by overlaying 2015 and 2004 building polygon data in order to track the persistency, demolition, construction and reconstruction of buildings over time. Since buildings are not identified across time our links rely on a shape matching algorithm. For each building, the algorithm determines whether it was there in the other period, or not, by comparing it with the buildings that overlap in the other time period.

This task is not straightforward, since the same building can be recorded in different ways depending on the aerial imagery used, whether building height was available, and the idiosyncrasies of the human digitizer.

Data and definitions

For 2004 we use the building dataset received from the Nairobi City Council with digitized polygons for every building, roughly 340,000 in the administrative boundary of Nairobi. For 2015 we use the dataset that was created by Ramani Geosystems using imagery (10-20cm resolution).

The nomenclature we use is as follows. First, a *trace* is the collection of polygon vertices that make up its outline. A *shape* is the area enclosed by the trace, and can be thought of as a representation of the rooftop of a building. A *cavity* is an empty hole completely enclosed in a shape. A *candidate pair* is the set of any two shapes in different time periods which spatially intersect. A *link* is the relationship between a set of candidates in one period to a set of candidates in the other time period.

Pre-processing

Before running our shape matching algorithm we clean up the data sets. First we take care of no data areas. There are some areas that were not delineated in 2004, including the Moi Air Base, and Nairobi State House. We drop all buildings in these areas for both 2004 and 2015, amounting to roughly 1,500 buildings from the 2015 data, and 100 buildings from 2004. Next we deal with overlapping shapes, an issue arising in the 2015 data, although not that for 2004. This is most often the same building traced multiple times. We identify all such overlapping polygons and discard the smaller version until no overlaps remain; about 1,400 buildings from the 2015 data this way. We also drop small shapes, in part because the 2015 data has many very small shapes, while the 2004 data does not. In order to avoid complications of censoring in the 2004 data, we simply drop all shapes that have an area of less than $1m^2$. We drop 2 small buildings in 2004, and 462 small buildings in 2015.

Another issue is that buildings are often defined as contiguous shapes in 2004, but broken up in 2015. For the majority of buildings we cannot aggregate the broken up pieces in 2015 since it is hard to identify such cases in general. To match these cases across time we rely on our one to many, and many to many matching algorithms defined below. However, in the specific case where a building is completely enclosed in another the task is much easier. First, we find all cavities present in each period, then we take all building shapes that overlap with the cavities in the same time period. After identifying all shapes that intersect a cavity, we redefine both shapes, the original shape containing the cavity and the shape intersecting it, as a single new shape.

Shape Matching Algorithm

After the pre-processing of each cross-section is complete, we run our shape matching algorithm to establish links between buildings across time periods. For any given building we consider 5 possible scenarios; that it has a link to no building, that it has a link to one building (one to one match), that it has a link to multiple buildings (one to many), that it is part of a group of buildings that match to one building (many to one), or that it is a part of a group of buildings that matches to a group of buildings (many to many). We follow and approach similar to Yeom et al (2015) however, due to the inherent difficulty of inconsistent tracings we contribute to their method by introducing the one to many and many to many approaches. We assign each link a measure of fit that we call the overlay ratio. We then choose optimal links based on the overlay ratio. Finally, we categorize links as matched or not using a strict cut-off on the overlay ratio of 0.5. Other cu-offs such as 0.4, 0.6 and 0.7 produced more errors in categorization.

Candidates

For all buildings A in the first time period, and B in the second time period we identify the set of candidates:

$$CP = \{(A, B);\ Area(A \cap B) \neq 0\}$$

For each candidate pair we find the ratio of the intersection area over the area of each shape, so if shapes A and B intersect, we find $r_{AB} = \frac{Area(A \cup B)}{Area(A)}$ and $r_{BA} = \frac{Area(A \cap B)}{Area(B)}$
We link all shapes which do not belong to a candidate pair to the empty set.

One to One Matching
First we consider candidate pairs to be links on their own. For each pair, we calculate the overlay ratio as the intersection area over union area, so if A and B are candidate pair, we find:

$$R_{AB} = \frac{Area(A \cap B)}{Area(A \cup B)} = \frac{Area(A \cap B)}{Area(A) + Area(B) - Area(A \cap B)}$$

One to Many Matching
For each time period separately, we identify all candidate pair links for which their intersection to area ratio is above threshold θ. For shape A we define a group $= \{B; r_{BA} \geq \theta\}$. Now we calculate the overlay ratio of one to many links as the intersection area over union area ratio:

$$R_{AG} = \frac{Area(A \cap \bigcup_{B \in G} B)}{Area(A \cup \bigcup_{B \in G} B)} = \frac{\sum_{B \in G} Area(A \cap B)}{\sum_{B \in G} Area(A \cup B)}$$

Many to Many Matching
Here we have two cases, one when the shapes are fairly similar, which we capture in previous sections (one to one, or many to one). The other is inconsistent shapes that form the same structure. To capture these we consider both time periods at the once, we clean the candidate pair list, keeping links for which either ratio is above a threshold θ_1:

$$LC = \{(A, B);\ r_{AB} \geq \theta_1 \ or \ r_{BA} \geq \theta_1\}$$

Then we condition to only keep shape for which the total ratio intersection is above threshold θ_2, so shape A will be included if $\sum_{B \in \{x | (A,x) \in LC\}} r_{AB} \geq \theta_2$. Now we are left with a new candidate list, which we convert to sets $LC = \{(\{A\}, \{B\})\}$ and start merging them:

$$if\ G_i \cap G_j \neq \emptyset \ or \ H_i \cap H_j \neq \emptyset : LC = \{(G_i \cup G_j, H_i \cup H_j)\} \cup LC / \{(G_i, H_i), (G_j, H_j)\}, \mathrm{i} \neq \mathrm{j}$$

We keep doing this until we can no longer merge any two rows. At this point we calculate the overlay ratio of many to many links as the intersection area over union section ratio:

$$R_{GH} = \frac{Area(\bigcup_{A \in G} A \cap \bigcup_{B \in H} B)}{Area(\bigcup_{A \in G} A \cup \bigcup_{B \in H} B)}$$

ICP Translation
We encounter a problem when the two shapes or groups of shapes are similar but do not overlap well, this usually stems from the angle at which the images were taken, and is especially prevalent with tall buildings. To address this issue, we translate one trace towards the other, and then recalculate the overlay ratio. As in Besl and McKay (1992), we use the iterative closest point (ICP) method to estimate this translation. To perform the ICP we ignore any cavity points as we found they often cause less suitable translation. We found that for similar shapes this will optimize the intersection area.

Optimal Linking
In the end, we rank all links by their overlay ratio. We iteratively keep the link with the highest overlay ratio, or discard it if at least one of the buildings in the link has already been confirmed in a separate link. From the list of optimal links, we define a link to be a match if its overlay ratio, or the overlay ratio after ICP translation is above 0.5. We then define all matched candidates as unchanged, and the remaining candidates as redeveloped. All buildings that were not considered as candidates are defined as infill, if from 2015, and demolished, if from 2004.

Accuracy Assessment

In order to assess the performance of the polygon matching algorithm we manually classified links between 2004 and 2015 for a random sample of buildings. We sampled 48 150x150m grid cells, stratifying over slum, non-slum within 3km, non-slum within 6km, and non-slum further than 6km to the CBD. The sample consists of over 2,250 buildings in 2004 and 3,500 buildings in 2015.

Results

We first break down matches by their mapping type. There are five types of manual link: redeveloped/infill/demolished (0), one to one match (1), one to many match (2), many to one match (3), and many to many match (4). For the algorithm we further split (0) into infill/demolished (-1) and redeveloped (0). Appendix table 1 shows the correspondence between the two mappings by building (a) and roof area (b). We can see that most errors come from the one to one matches, however, the many to many matches have the worst performance. Overall the diagonal values are quite high, which means not only are we matching buildings well, but also the algorithm is recognising the clumping of buildings as a human does (bear in mind that, for example, the one to one matches which we 'misclassify' as many to many will still be classified as match in the final data). Finally, we have perfect correspondence for demolition and in 2015 nearly perfect for infill.

Next we compare buildings that were matched by the algorithm and those matched manually. For now we use a cut-off of the overlay ratio of 0.5, later we explore the effect of different cut-offs on performance. As seen in appendix table 1 infill and demolition are classified with almost perfect correspondence. For this reason we ignore buildings with these mappings and focus on accuracy of redevelopment and unchanged. In appendix table 2 we condense mappings 1, 2, 3, and 4 into category 1, while redevelopment, or category 0, remains the same.

We define precision *P* (negative predictive value *NPV*) as the fraction of buildings classified as unchanged (redeveloped) by the algorithm that are correct, recall *R* (true negative rate *TNR*) as the fraction of buildings classified as unchanged (redeveloped) by hand that the algorithm gets correct, and the F1 score (F) as the weighted average of the two.

$$P = \frac{True\ Positive}{Positive\ Predictions}, \qquad NPV = \frac{True\ Negative}{Negative\ Predictions}, \qquad R = \frac{True\ Positive}{Positive\ Condition},$$

$$TNR = \frac{True\ Negative}{Negative\ Condition}, \qquad F = \frac{2 * P * R}{P + R}$$

The confusion matrix in table 2 is done across all sampled buildings in 2004 and weights observations by buildings (1) and roof area (2). The F1 score is high in both cases, but in part this is due to relative success classifying unchanged buildings: precision for buildings that were classified as redeveloped by the algorithm is 76% of buildings and 72% of roof area, while recall of true redeveloped buildings is 83% of buildings and 74% of roof area

In our first attempt we arbitrarily picked 50% as a cut off of the overlay ratio. Here we take a closer look at this choice. Using our manually classified links we can maximize the F1 score with respect to the cut off. In appendix figure 1 we plot the F1 score weighted by roof area against cut-offs of the overlay ratio for the 2004 data. We find that the highest F1 score comes just below 50% suggesting our first estimate was not far off.

In figure 1 we plot lines for each method of calculating the overlay ratio: without ICP, with ICP, and the maximum of the two. Around 50% we can see that the maximum performs best, but with only a very slight improvement over the ICP alone, which is in turn marginally better than without the ICP.

Appendix Table 2.1 – Mapping Correspondence 2004

a) Weighted by Building						
	Algo=-1	Algo=0	Algo=1	Algo=2	Algo=3	Algo=4
Manual=0	280	433	41	16	11	20
Manual=1	0	25	712	10	1	25
Manual=2	0	29	21	266	0	20
Manual=3	0	18	6	0	137	1
Manual=4	0	65	52	24	63	135
b) Weighted by Area (sq-m)						
	Algo=-1	Algo=0	Algo=1	Algo=2	Algo=3	Algo=4
Manual=0	12708	28187	4913	2780	943	1043
Manual=1	0	908	112762	4180	279	1775
Manual=2	0	3575	2328	89472	0	2819
Manual=3	0	910	1053	0	14148	23
Manual=4	0	5317	5528	4795	4464	14262
Mapping definitions: -1 demolition or infill; 0 redevelopment; 1 one to one match; 2 one to many match; 3 many to one match; 4 many to many match						

Appendix Table 2.2 – Matching all areas 2004

a) Weighted by Building			
	Algo=0	Algo=1	Recall
Manual=0	433	88	0.83
Manual=1	137	1473	0.91
Precision	0.76	0.94	F=0.93
b) Weighted by Area (sq-m)			
	Algo=0	Algo=1	Recall
Manual=0	28187	9679	0.74
Manual=1	10710	257888	0.96
Precision	0.72	0.96	F=0.96

Appendix Figure 2.1

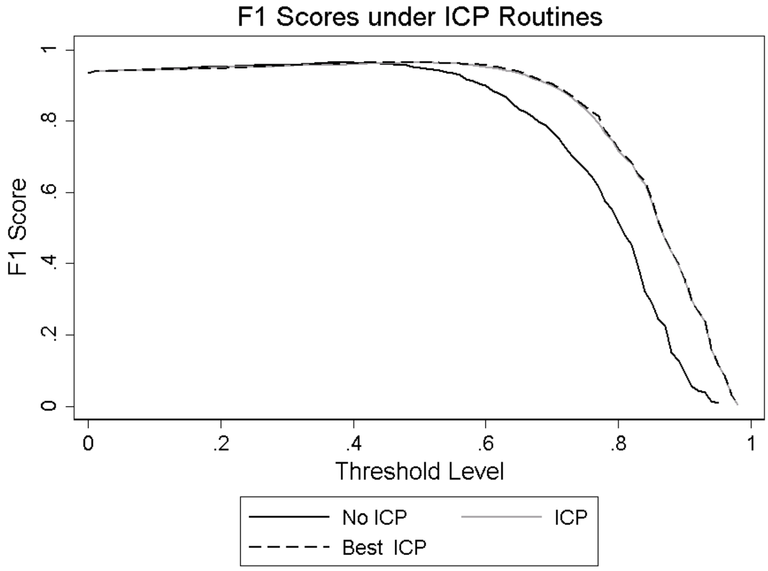

Appendix Table 2.3: List of public uses	
Recreational a) Impala club, Kenya Harlequins, and Rugby Union of East Africa (0.14kmsq) b) Golf Course (0.9kmsq) c) Arboretum (0.25kmsq) d) Central park, Uhuru park, railway club, railway golf course (0.5kmsq) e) Nyayo stadium (0.1kmsq) f) City park, Simba Union, Premier Club (1.1kmsq) g) Barclays, Stima, KCB, Ruaraka, Utali clubs, and FOX drive in cinema (0.3kmsq) Undeveloped a) Makdara Railway Yard (1kmsq) b) John Michuki Memorial Park (0.1kmsq) Special use -- Includes poorly traced areas a) State House b) Ministry of State for Defence c) Forces Memorial Hospital and Administration Police Camp d) Langata Army Barracks e) Armed Forces f) Moi Airbase	Public utility a) Dandora dump (0.5kmsq) b) Sewage works (0.25kmsq) g) Kahawa Garrison Public use a) Communications Commission of Kenya (0.1kmsq) b) Langata Womens prison (0.2kmsq) c) Nairobi and Kenyatta hospitals, Milimani Police Station, Civil Service club d) Mbagathi hospital, Kenya Medical Research Institute, Monalisa funeral home e) National museums of Kenya f) Kenya convention centre and railway museum g) Industrial area prison h) Mathari mental hospital, Mathare police station, traffic police, Kenya police, Ruaraka complex, and National youth service i) Jamahuri show ground Educational (not primary and secondary schools) a) University of Nairobi and other colleges b) Kenya Institute of Highways & Built Technology c) Railway Training Institute d) Kenya Veterinary Vaccines Production Institute e) Moi Forces Academy f) NYS engineering, Kenya Institute of Monetary Studies, KCA university, KPLC training, Utali college

Appendix References

Besl, P., and Mckay, N. 1992. “A Method for Registration of 3-D Shapes.” IEEE Transactions on Pattern Analysis and Machine Intelligence. 14(2), 239—256

Google Earth V 7.3.0.3830. (February 2, 2002). Nairobi, Kenya. 1° 18’ 19.09”S, 36° 49’ 16.26”W, Eye alt 4.61km. DigitalGlobe 2017. http://www.earth.google.com [August 18, 2017].

Google Earth V 7.3.0.3830. (December 15, 2003). Nairobi, Kenya. 1° 18’ 19.09”S, 36° 49’ 16.26”W, Eye alt 4.61km. DigitalGlobe 2017. http://www.earth.google.com [August 18, 2017].

Google Earth V 7.3.0.3830. (February 14, 2004). Nairobi, Kenya. 1° 18’ 19.09”S, 36° 49’ 16.26”W, Eye alt 4.61km. DigitalGlobe 2017. http://www.earth.google.com [August 18, 2017].

IPE Global Private Limited and Silverwind Consultants, 2013. “Consultancy Services for City/Municipal Situational Analysis of Conditions of Informal Settlements in 15 Municipalities.” Prepared for the Government of Kenya, Ministry of Lands, Housing and Development.

Sarah Williams, Elizabeth Marcello & Jacqueline M. Klopp, 2014. “Toward Open Source Kenya: Creating and Sharing a GIS Database of Nairobi,” Annals of the Association of American Geographers, 104:1, 114-130.

USGS (2005), Shuttle Radar Topography Mission, 1 Arc Second scene SRTM_f03_s002e036, Fnished Version 3.0, Global Land Cover Facility, University of Maryland, College Park, Maryland.

Yeom, J., Kim, Y., and J. Lee, 2015. “Hierarchical ICP Matching to Enable Interoperability of Multisource Building GIS Data.” Advanced Science and Technology Letters, Vol. 89 (Architecture and Civil Engineering) pp. 55-60.

Zinnes, C., et al., 2012. Kenya Urban Program Baseline Study (NORC, University of Chicago, Chicago); www.norc.org/Research/Projects/Pages/kenya-urban-program-baseline-study.aspx