

WIDER Working Paper 2019/111

The effects of the Maputo ring road on the quantity and quality of nearby housing

Peter Fisker,¹ David Malmgren-Hansen,² and Thomas Pave Sohnesen³

December 2019

Abstract: Using convolutional neural networks applied to satellite images covering a 25 km x 12 km rectangle on the northern outskirts of Greater Maputo, we detect and classify buildings from 2010 and 2018 in order to compare the development in quantity and quality of buildings from before and after construction of a major section of ring road. In addition, we analyse how the effects vary by distance to the road and conclude that the area has seen large overall growth in both quantity and quality of housing, but it is not possible to distinguish growth close to the road from general urban growth. Finally, the paper contributes methodologically to a growing strand of literature focused on combining machine-learning image recognition and the availability of high-resolution satellite images. We examine the extent to which it is possible to exploit these methods to analyse changes over time and thus provide an alternative (or complement) to traditional impact analyses.

Key words: Mozambique, infrastructure, housing convolutional neural networks, remote sensing, impact assessment

JEL classification: O18, R14

Acknowledgements: The authors acknowledge comments and suggestions from participants at the Annual Conference of the Inclusive Growth in Mozambique project held in Maputo in November 2019.

¹Development Economics Research Group (DERG), Department of Economics, University of Copenhagen, Copenhagen, Denmark, corresponding author email: pkf@econ.ku.dk. ²Technical University of Denmark (DTU), Copenhagen, Denmark., ³Independent Consultant, Copenhagen, Denmark.

This study has been prepared within the UNU-WIDER project [Inclusive growth in Mozambique – scaling-up research and capacity](#).

Copyright © UNU-WIDER 2019

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9256-747-7

<https://doi.org/10.35188/UNU-WIDER/2019/747-7>

Typescript prepared by Lesley Ellen.

The United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland, Sweden, and the United Kingdom as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the Institute or the United Nations University, nor the programme/project donors.

1 Introduction

In 2011, the Government of Mozambique launched the Maputo Ring Road construction project, which aimed to connect the main entry roads into the city and allow traffic to and from South Africa to bypass the city centre. Work began in 2012 and most of the road was finished by 2015, although one intersection (Tchumene) was still under construction at the time of writing but is expected to be completed early in 2020. The total cost of the project is around US\$315 million, of which US\$300 million has been financed through a loan from the Chinese state-owned bank Exim-Bank. The project has included improvements to 22 km of existing road and the construction of 52 km of new road (Government of Mozambique 2012). The latter will be the main focus of this paper.

The economic benefits of the ring road—mostly in the form of shorter travel times—have already spread throughout Maputo and most of southern Mozambique. Locals enjoy a shorter commute, business owners have lower transportation costs, and tourists can more easily reach the sandy beaches of Mozambique.¹ Further, following completion in 2017 of the Maputo Bay Bridge (the largest suspension bridge in Africa), which links Maputo to its southern hinterland and to Durban, the road infrastructure in Mozambique’s capital has had a much-needed upgrade, relieving pressure on several transport bottlenecks. This study will not attempt to assess the impacts of these infrastructure projects on the overall economy of Mozambique or Maputo. Its main purpose is to evaluate the changes to housing conditions in the areas surrounding the newly built sections of the ring road.

When observing aerial photography or high-resolution satellite images of the northern parts of Greater Maputo, it is clear that widespread urbanization has occurred during the past decade. Some of this may have happened irrespectively of the construction of the ring road; however, academics and policy makers have an interest in knowing more about how the ring road project has contributed to this development. For instance: Where specifically has urban expansion taken place? Are the developments concentrated around the newly built stretch of road, or are they determined by distance to the city centre? Can we learn anything about the quality of houses in this ‘new’ part of the city, and how does proximity to the road affect this, if at all? Knowing the answers to these questions and more can shed light on some of the effects of improving urban infrastructure which are often not measured by interview-based impact assessments.

Using convolutional neural networks (CNNs) trained on around 200,000 tagged satellite images of structures and non-structures in urban areas in Mozambique, we detect and classify buildings in high-resolution satellite images for the years 2010 and 2018 in order to compare the distribution of houses before and after construction of the road. In addition, we estimate changes to the share of buildings with high-quality roofs and the share of buildings under construction, and analyse the effect on these estimates arising from their distance to the road. Finally, the paper contributes to a growing strand of literature focused on combining recent advances in machine-learning image recognition and the availability of high-resolution satellite images. In particular, we test the extent to which it is possible to exploit these methods to analyse changes over time and thus provide an alternative (or complement) to traditional impact analyses.

¹ According to one participant at the Annual Conference of the ‘Inclusive Growth in Mozambique’ project held in Maputo in November 2019, the ring road has reduced his daily commute from 120 minutes to 20 minutes.

1.1 Context and literature

The availability of high-resolution images of the earth is increasing rapidly, leading to interesting new applications. The availability of these images also provides opportunities to improve public policies, especially in data-scarce environments, including many developing countries. While the increase in the number of satellite images, such as the free images provided by the European Space Agency's sentinel satellites, is partly due to government investment, it may be due more to the expansion of commercial satellites. The imaging company Planet Labs currently maintains 140 satellites, enough to pass over every place on earth once a day. Maxar, formerly DigitalGlobe, is building a constellation that will be able to revisit spots 15 times a day, and BlackSky Global promises to revisit most major cities up to 70 times a day (Beam 2019).

The usability of satellite images for public policies has improved markedly through the application of CNNs for image recognition. For instance, CNNs can detect specific features in images which can support analysis and policies. Examples include identification of washed-away buildings after a tsunami (Fujita et al. 2017); detection of building footprints, including in many African countries (Tiecke et al. 2017; Liu et al. 2019); tracking of icebergs at sea (Malmgren-Hansen et al. 2019), 24/7 surveillance of the quality of infrastructure such as bridges (Avci et al. 2017); and estimation of poverty at the neighbourhood level in urban areas in Mozambique (Fisker et al. 2019a). This study builds upon partial results from Fisker et al. (2019a), namely the identification of structures of various types and qualities from high-resolution satellite images.

Lall et al. (2017) set the context for describing the potential of infrastructure investments in African urban centres by listing a number of stylized facts about large African cities that constrain them from reaching their growth potential. First, they tend to be crowded with people rather than physical capital. Central business districts are generally characterized by a high share of paved roads, which is comparable to same-size cities in the Western world. However, paved roads are very scarce outside of the city centres. This is not optimal since it generates congestion problems in the centre as well as long commutes for people living in the periphery. Second, cities are made up of separate neighbourhoods that are not well connected. This means that the number of available jobs within a reasonable travel distance from a given location is low. This is especially problematic for the formal sector, as informal jobs are often carried out in or close to people's homes. Third, and perhaps as a result of the lack of physical infrastructure, population density is also high close to the city centres, but quickly tapers off in the outskirts. In Maputo, 30 per cent of land within 5 km of the central business district is left undeveloped (Lall et al. 2017).

In one of the few studies which attempts to measure the citywide effects of the construction of new transport routes, Bernard et al. (2016) show that, in Kampala, a northern bypass (similar in size to the Maputo ring road) has reduced the cost of living for high-skilled labour by 19 per cent and by 3 per cent for low-skilled labour—a welfare effect of eight times the construction costs. Their model also captures the direct impact on transport costs as well as the more long-term agglomeration effects that arise from the location choices of firms and households.

1.2 Hypotheses

The findings of the previous studies described above lead to the formulation of the following hypotheses about the potential effects of the construction of the Maputo ring road:

- i) As lower costs of living and shorter travel time should attract households and firms, we expect the structure density to increase over time, and to a greater extent in the vicinity of the road than in control areas.

- ii) Completion of the ring road in Maputo should lead to higher welfare citywide (not tested here), and especially in areas around the road. It is therefore expected that housing quality—as measured by the share of high-quality roofs within a certain geographical area—increases after completion of the road, and to a greater extent in the vicinity of the road than in control areas.

2 Data generation

The area of interest studied is illustrated in Figure 1, where the left side illustrates the various stages of the ring road project with an ellipsoid highlighting the areas of interest in this study. The exact study area is defined in the right-hand panel, a rectangle covering about 25 km x 12 km spanning the northern outskirts of the city of Maputo and the large areas of the Matola and Marracuene suburbs.

Figure 1: Study area



Sources: Wikipedia (2019) on the left, and Google Earth on the right.

In order to track the city’s development, two CNNs were trained to detect 1) the presence of structures, and 2) the type or quality of the detected structures. The CNNs were first developed for the purpose of generating detailed urban poverty maps for urban areas in Mozambique (Fisker et al. 2019a). With minor adjustments, the same algorithms are applied to images before and after the construction of the ring road, and thereby provide evidence of the impact of the road on city development.

2.1 Developing CNNs for structure detection

The CNN’s ability to recognize structures in images is based on training data. Hence, the CNN is trained by showing it a large number of examples of the object it has to recognize. This paper relies on two separate CNNs: 1) the CNN-detector, which detects structures, and 2) the CNN-classifier, which subsequently categorizes the detected structures from the CNN-detector into different types. The CNN algorithms were trained using Python programming language with Tensorflow and Keras libraries. For both CNNs, the training images are 50 x 50 pixels, which is big enough to cover a typical house and its surroundings. The images for training the CNNs were downloaded from Google Maps API, under Google’s ‘fair use’ agreement, making them free for non-commercial use. The training data is publicly available through Mendeley Data (Fisker et al. 2019b).

The dataset for training the CNN-detector consisted of a total of 107,507 tagged structures from the five largest cities in Mozambique (Maputo, Nampula, Tete, Quelimane, and Beira). Further, tags indicating where there were no structures were also required to train the CNN-detector. For this, 100,000 coordinates, at least 20 metres from the nearest house tag in the dataset, were extracted for the city of Tete. Tete was chosen as it was completely tagged, unlike the other cities, which were only partly tagged (Fisker et al. 2019a).

For the second-stage CNN-classifier, the detected structures were classified into ‘apartment’, ‘non-residential’, ‘under construction’, ‘painted roof’, ‘grey roof’, and ‘palm roof’ categories. After reviewing the data, the ‘apartment’ and ‘palm roof’ categories were found to be unsuitable for CNN training due to under-representation of samples or the tagging quality. For the training of the CNN-classifier, a balanced sample was extracted with approximately equal samples of the ‘non-residential’, ‘under construction’, ‘painted roof’, and ‘grey roof’ categories. In this study, we only employ the CNN-detector and the CNN-classifier ‘painted roof’ and ‘under construction’ categories.

An out-of-sample evaluation of the performance of both the CNN-detector and CNN-classifier found that the CNN-detector has an accuracy of 97 per cent (evaluated on 10 per cent of training sample excluded) whereas the CNN-classifier has, on average, 67 per cent accuracy (painted roof: 73 per cent, under construction: 69 per cent) (Fisker et al. 2019a).

2.2 CNNs applied to the ring road

To analyse the effects of the construction of the Maputo ring road on local housing, we applied the observed relationship between the images and the structure categories obtained in the previous step to new images of the study area depicted in Figure 1.

In order to ensure 1) the highest possible resolution, 2) no gaps or overlaps between images, and 3) minimal manual work, a so-called fishnet grid of 220 cells of around 1,155 m x 1,155 m was created using GIS software. An example of one of these locations for 2010 and 2018 is shown in Figure 2. The somewhat odd cell size corresponds exactly to the area of the images that are downloadable from Google Earth when the resolution is set to the maximum 4,000 x 4,000 pixels at the highest meaningful zoom level. The centre points of the fishnet grid are then exported to KML format and opened in Google Earth, making it possible to download each image with only a few clicks. The images used in this study had to be manually downloaded because, at the time of writing, Google Maps API only allowed automatized extraction of images by centre coordinates for the most recent images in their archive. This would not have allowed us to compare images going back in time.

From these images, centre coordinates of sub-patches detected as structures were stored and converted to latitude and longitude by approximation from the centre of the cells. We also calculated mean statistics for the whole cells, such as construction density, given as the number of sub-patches detected as containing a structure, divided by the total number of sub-patches in the cell.

Based on the extracted coordinates of each detection, we then created a dataset consisting of 25,252 ‘pixels’ of around 115 m x 115 m, as depicted in Figures 4, 5, and 6. Each pixel contains information on construction density (i.e. share of pixels covered by structures), share of high-quality roofs, share of buildings under construction, distance to city centre, distance to ring road, and whether the pixel is located north or south of the ring road.

Figure 2: Example of images from 2010 (left) and 2018 (right) of areas through which the ring road now runs.



Source: Google Earth.

Importantly, the 2010 baseline data was taken around a year before the Government of Mozambique decided to implement the construction project. We therefore consider it likely that the population of the general area were not expecting a road to be built. At the other end of the timescale, in 2018, at least three years had passed since the end of the road construction process (with the exception of the Tchumene intersection), meaning that most of the households and firms that wanted to move to the area had had enough time to realize their intentions. So, while the process goes on and new areas are probably still being populated, most of the ring road's immediate impact on housing can be assumed to have materialized in 2018.

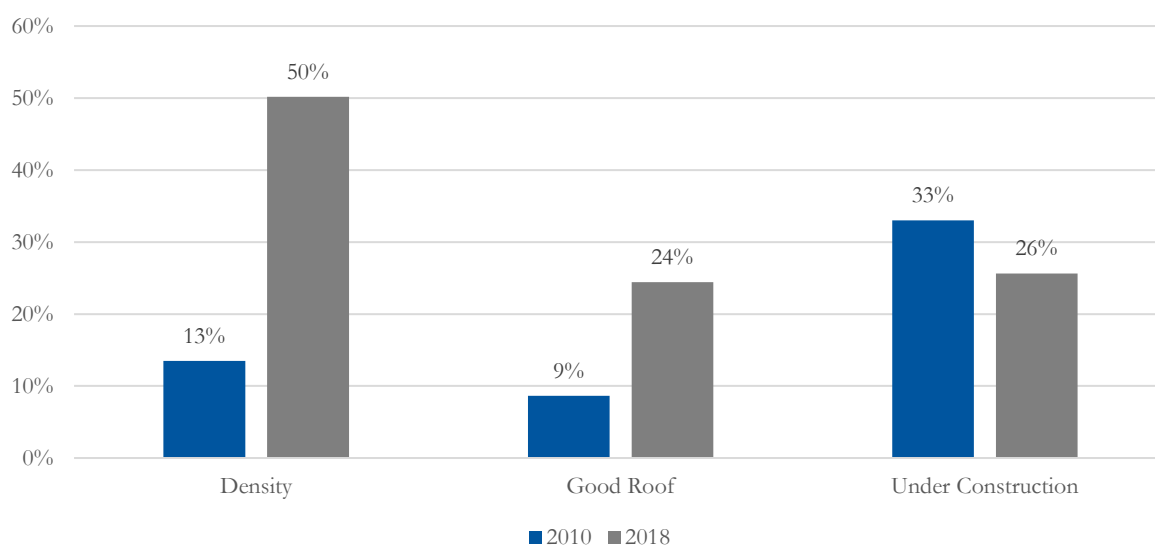
3 Analysis

Our analysis, based on the data generated by the CNNs, consists of two parts: a descriptive analysis and a regression analysis.

3.1 Descriptive analysis

The overall trends in our key variables—density of structures, share of high-quality roofs, and share of buildings under construction—are shown in Figure 3. It shows that construction density across the whole study area increased from 13 per cent in 2010 to 50 per cent in 2018, the share of high-quality roofs increased from 9 per cent to 24 per cent, while the share of buildings under construction decreased from 33 per cent to 26 per cent between 2010 and 2018. Since the last two are shares of the first, this means that the *density* of high-quality roofs increased far more dramatically than the 15 percentage point increase, and that, while the share of structures under construction decreased, the absolute number of structures under construction increased from 2010 to 2018.

Figure 3: Key variables in 2010 and 2018, per cent



Notes: n = 25,252. Density is the share of built-up area among pixels; Good Roof and Under Construction are shares of detections classified as such.

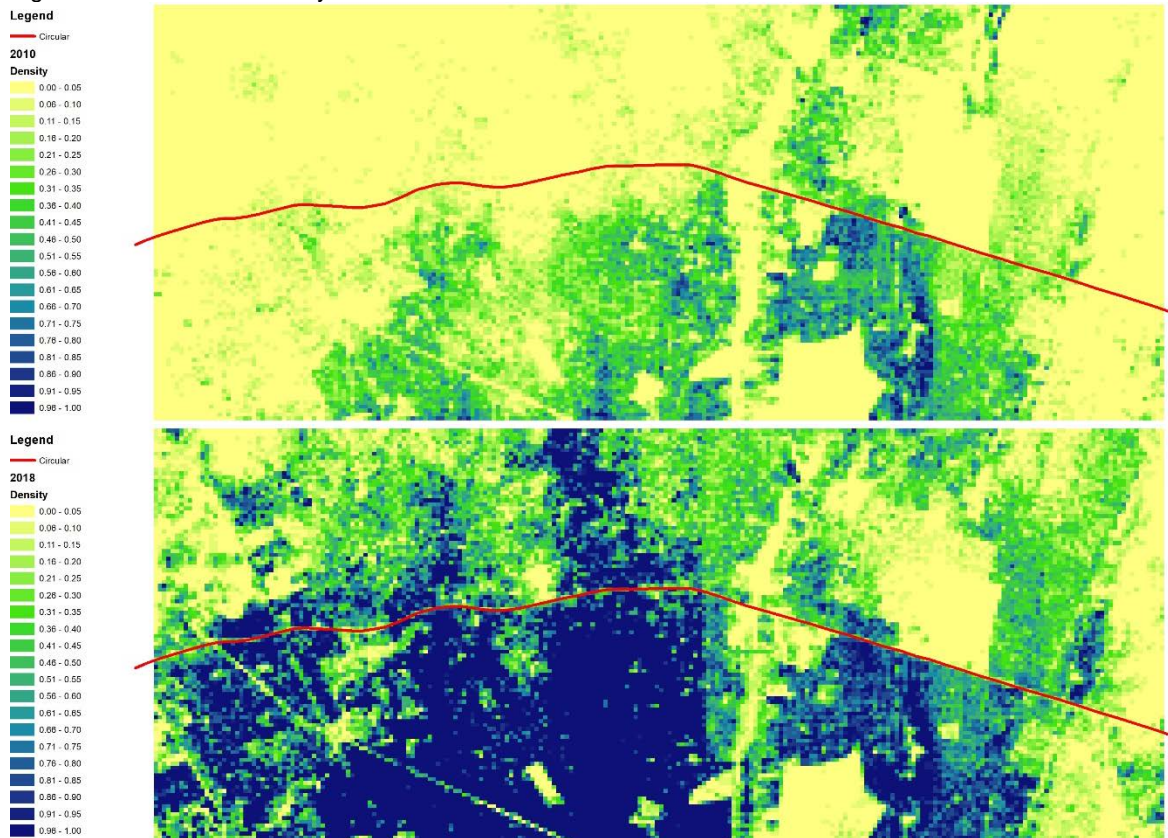
Source: authors' own calculations.

Figures 4, 5, and 6 illustrate the spatial distribution of the changes in structures over time. Figure 6 shows that, in 2010, most construction activity took place in the southern area, closer to the city centre, whereas the construction activity in 2018 was more concentrated in the north and central areas of the study area. The most dramatic development between 2000 and 2018 in terms of housing quantity and quality took place in the western half of the study area (Figures 4 and 5). This area officially forms part of the suburb of Matola, which is in a different administrative division to Maputo City, namely the Province of Maputo and Municipality of Matola. However, areas located in the south-east quadrant (and thereby part of Maputo City) were already partly built-up in 2010. The north-eastern quadrant of the maps—belonging to the Province of Maputo and the district of Marracuene—also do not show as strong growth in construction density as the western half. We have no evidence that the administrative boundaries are key to this, nor do we have any evidence that they are not important for this observed growth.

In Figures 4, 5, and 6, some areas are blank for both years studied. According to Andersen et al. (2015), these are military zones (bottom), state forests (just north of the road in the eastern half), or areas designated for urban agricultural production (far-east side, north–south central stretch).

It is noteworthy for all three city growth indicators (density of structures, share of high-quality roofs, and share of structures under construction) that the road is located in areas of great change, potentially indicating that the road was important for the observed growth. However, while the road is located in areas that have grown, there is also limited indication that the growth is centred around the road itself. None of Figures 4, 5, or 6 clearly show increased structure density, construction activity, or high-quality roofs in the immediate vicinity of the road. As explained earlier, the road is so influential for the city that it is not possible to identify areas that have not been impacted by it. We can only assess whether the road has had a larger impact on areas closer to it than on those further away. To examine this in more detail we turn to our regression analysis.

Figure 4: Construction density in 2010 and 2018



Notes: image based on 25,252 pixels and detection of structures by the CNN-detector model.

Source: authors' own calculations.

Figure 5: Density of high-quality roofs in 2010 and 2018

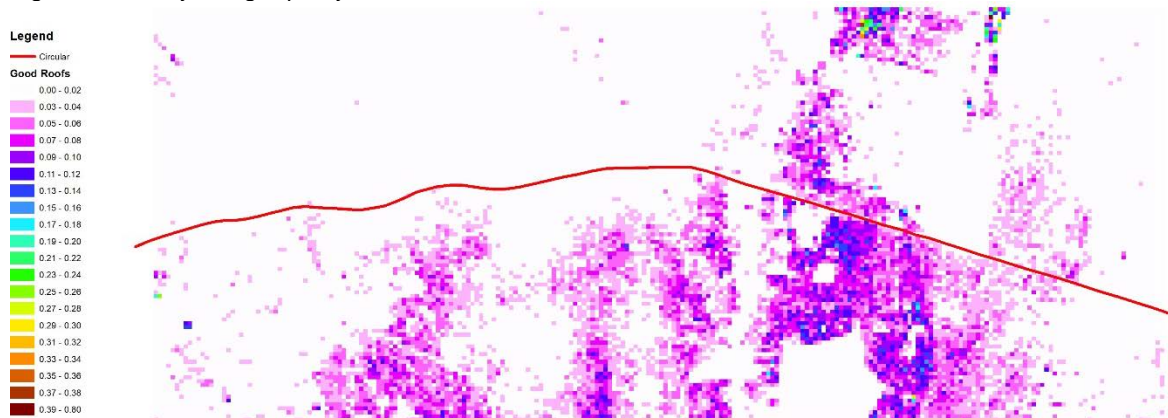
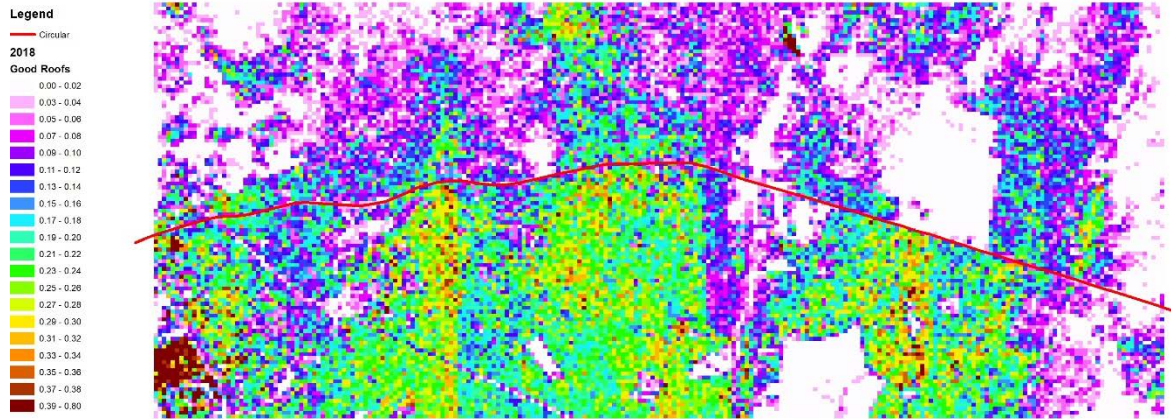


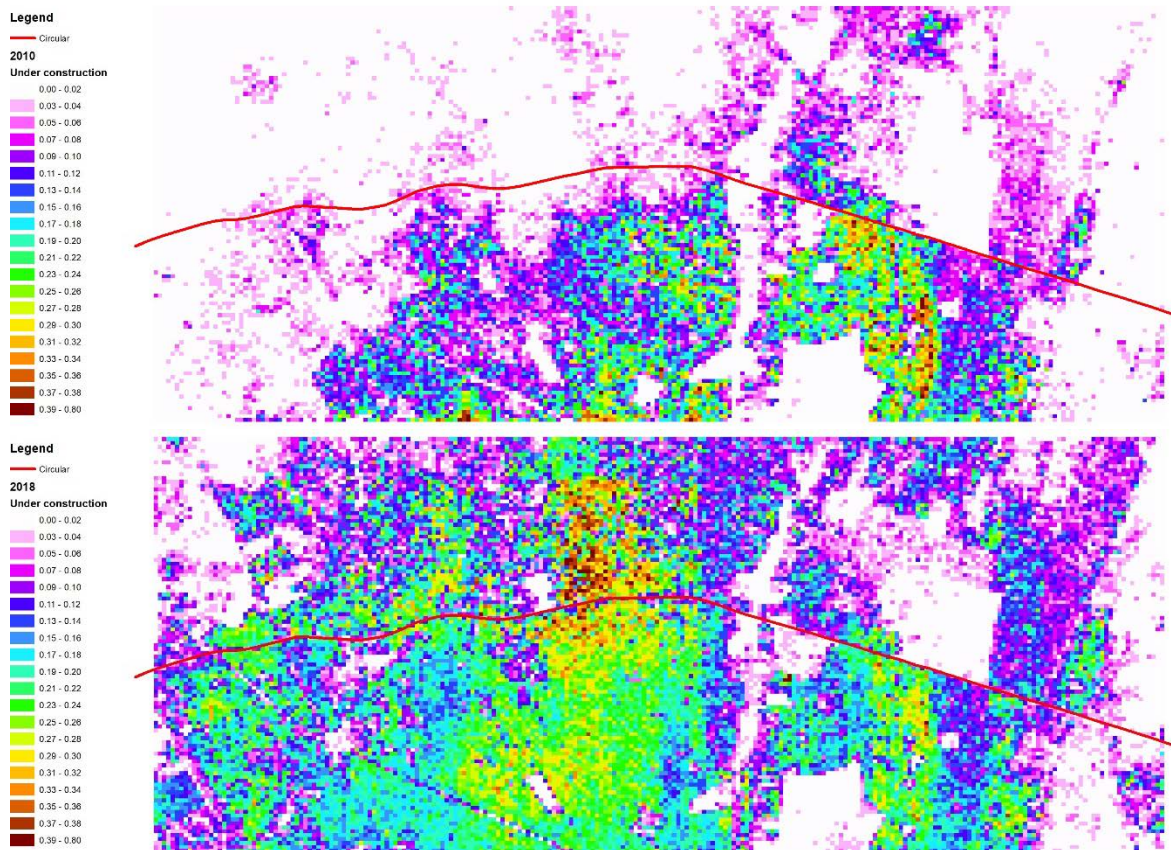
Figure 5 (cont.): Density of high-quality roofs in 2010 and 2018



Notes: image based on 25,252 pixels and detection of high-quality roofs by the CNN-classifier model.

Source: authors' own calculations.

Figure 6: Density of buildings under construction in 2010 and 2018



Notes: image based on 25,252 pixels and detection of type under construction by the CNN-classifier model.

Source: authors' own calculations.

3.2 Regression analysis

To assess the road's impact, we focus on the change in construction density, construction and quality of roofs between 2010 and 2018, as in regression (1):

$$\Delta Y_i = \beta_0 + \beta_1 * dist_road_i + \beta_2 * dist_cc_i + \epsilon_i \quad (1)$$

where ΔY_i is the change in structure density, density of high-quality roofs, or density of buildings under construction between 2010 and 2018 in pixel i . dist_cc and dist_road measure the geodesic distance to the city centre and ring road in km respectively. β_0 is a constant and ϵ_i is a pixel-specific random error term. The key variable of interest is dist_road and β_1 , which captures the change over time in different distances to the road. The distance to the city centre is included as a control variable as we expect this to have a direct influence on the outcome variables beyond and irrespective of the presence of the ring road.

Implementing regression (1), with and without controlling for distance to the city centre, shows that the closer one gets to the road, the higher is the growth in structure density and high-quality roofs, but there are fewer structures under construction (interpretation is reverse of coefficients as our key variable measures distance to road and not proximity to road). This could indicate a positive impact on the density of structures and high-quality buildings close to the road that have already matured, and therefore fewer structures under construction.

Table 1: Growth in construction density, high-quality roofs, and under construction 2010–18

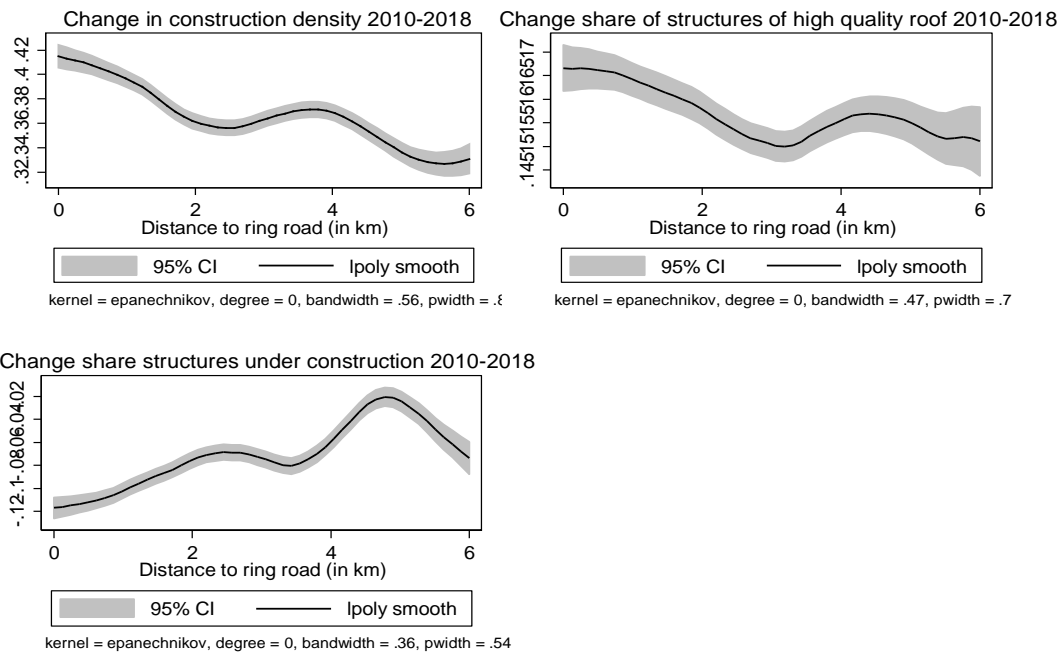
	Structure density		Share high-quality roof		Share under construction	
	(1) ΔY	(2) ΔY	(3) ΔY	(4) ΔY	(5) ΔY	(6) ΔY
Distance to road (km)	-0.015*** (0.001)	-0.015*** (0.001)	-0.002*** (0.001)	-0.002*** (0.001)	0.013*** (0.001)	0.012*** (0.001)
Distance to city centre (km)		0.001 (0.001)		-0.001** (0.000)		0.016*** (0.000)
Constant	0.409*** (0.004)	0.398*** (0.011)	0.163*** (0.002)	0.175*** (0.005)	-0.111*** (0.003)	-0.412*** (0.009)
N	25,252	25,252	25,252	25,252	25,252	25,252
r ²	0.01	0.01	0.00	0.00	0.01	0.05

Note: standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Source: authors' own calculations.

However, the impact of the road may also vary by the distance to the road. It might be preferable to be close to the road for easy access, but not be right next to it, as this could cause discomfort in terms of e.g. noise and heavy traffic. Figure 7 explores this by estimating the non-linear relationship between changes in our outcome variables of interest and the distance to the road. For growth in density of structure and share of high-quality structures, the relationship is largely linear. For the share under construction, there is a non-linear relationship with a peak in construction growth about 5 km from the road, which could be consistent with the current boundary of the city (Figure 6).

Figure 7: Non-linear relationship between outcomes and distance to road



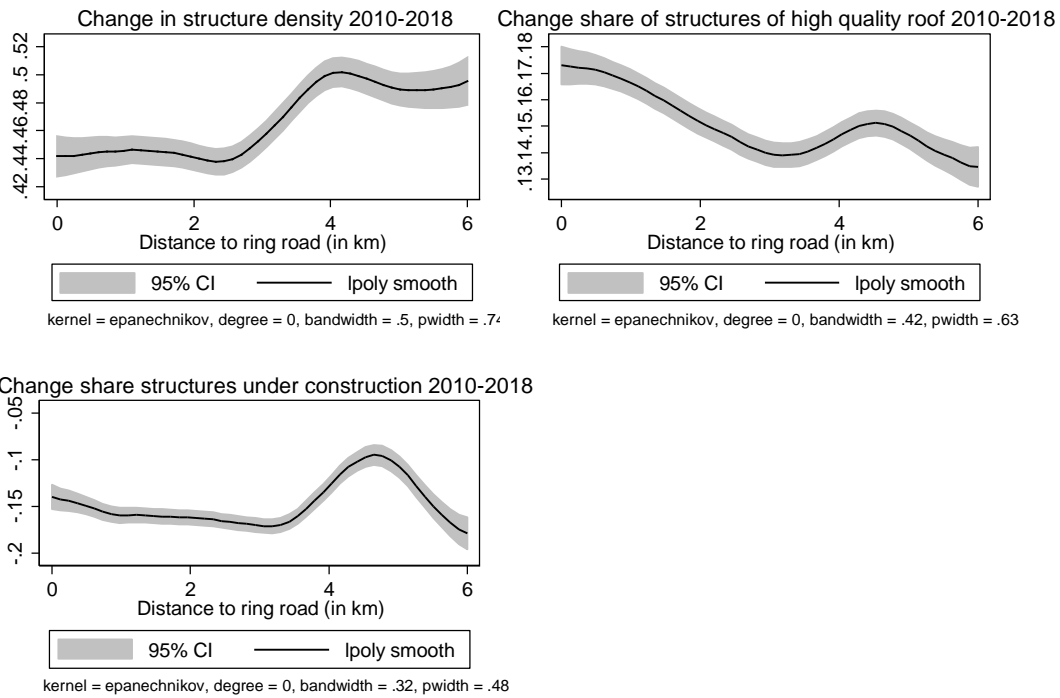
Source: authors' own calculations.

As highlighted in our descriptive analysis, in the study area, the city is growing rapidly from south to north and the ring road is right in the middle. Though growth in the variables of interest is higher in locations closer to the road than further away, the analysis presented thus far is not able to argue that the road—and not just the general growth of the city from south to north—is driving this. One way to argue that growth is being driven by the road, not just by general city expansion, is by analysing the impacts of the road north and south of the city separately. Figures 8 and 9 replicate Figure 7 for pixels located south and north of the ring road, respectively. The figures reveal that the impact of the road is very different in the north compared to the south of the road. The overall results seen in Figure 7 and Table 1 are clearly driven by different subsamples.

The higher growth in structure density and share under construction close to the road is driven by areas north of the road, while the opposite pattern is observed south of the ring road. The observed higher growth in high-quality housing closer to the road, on the other hand, is driven by areas south of the road, while little growth is observed for areas north of the road.

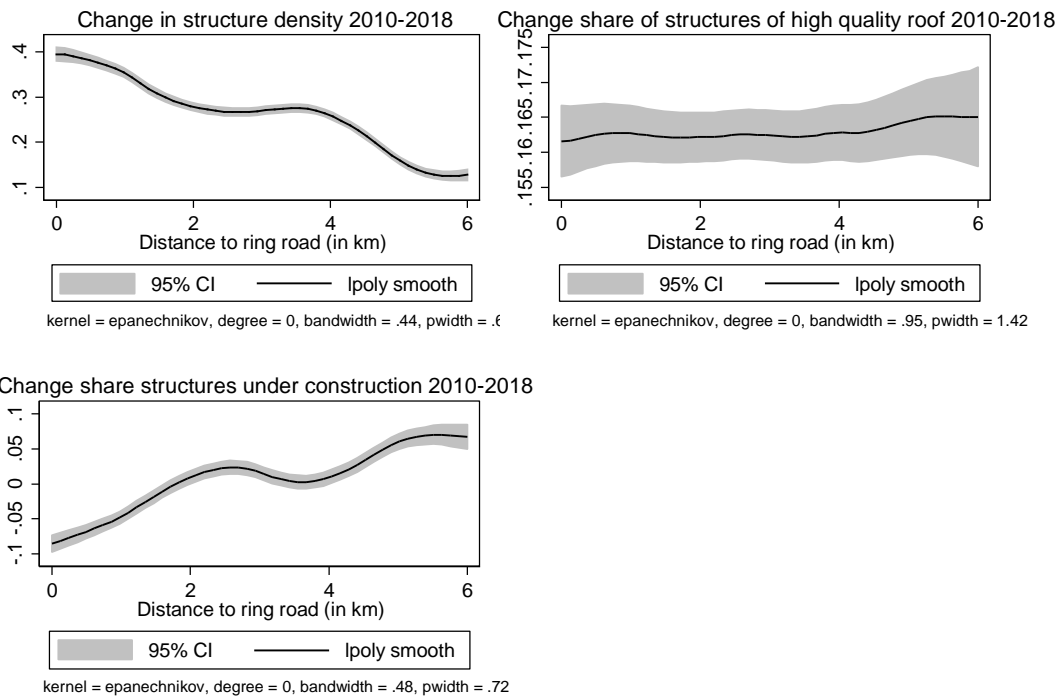
All in all, this points to the road being located in an area that is growing, but there is no indication of any increased growth in the density of structure, quality of structures, or structures under construction in the vicinity of the road that is in addition to the general growth of this part of the city.

Figure 8: Non-linear relationship between outcomes and distance to road, south of the road



Source: authors' own calculations.

Figure 9: Non-linear relationship between outcomes and distance to road, north of the road



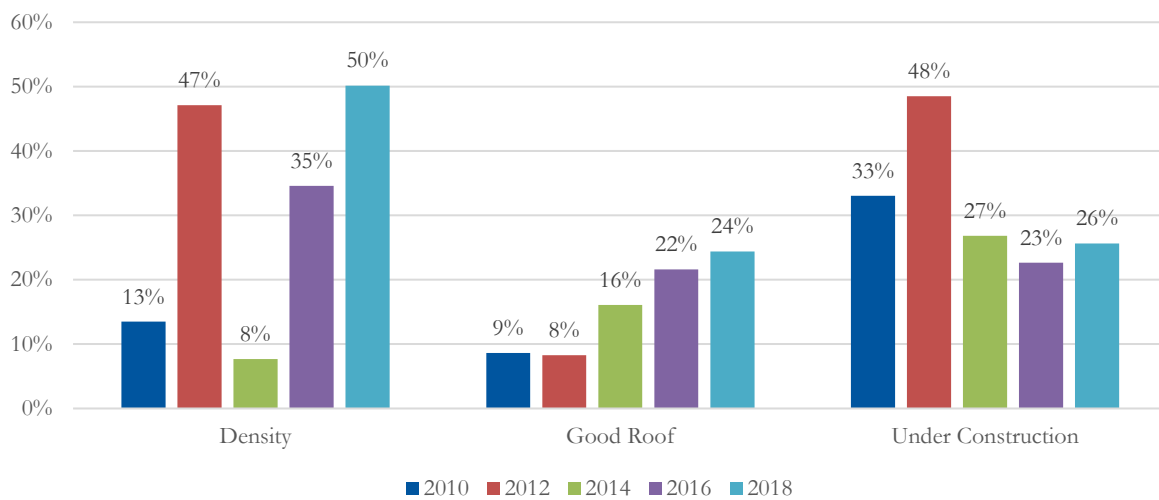
Source: authors' own calculations.

4 Challenges and caveats

We chose to compare images from the years 2010 and 2018 for two reasons. First, these two end years provide a good period for analysis since construction of the ring road commenced shortly after the baseline year of 2010, and 2018 has enough of a margin from the construction for the expected effects to have had time to materialize. Second, good-quality, cloud-free images were available for the two years, taken at around the same time of year, i.e. September for 2010 and May/June for 2018. While not exactly the same months, both periods are in the dry season.

In order to elaborate on the analysis, three extra sets of images were downloaded for the intermediate years of 2012, 2014, and 2016. However, these were not all of the same quality, and variation in the month and the time of day when the images were taken seem to have influenced the number of structures detected by the CNN structure algorithm in particular. This is apparent in Figure 10, which shows an impossible trend over time.

Figure 10: Key variables in 2010, 2012, 2014, 2016, and 2018.



Notes: $n = 25,252$. Density is the share of built-up area among pixels; Good Roof and Under Construction are shares of detections classified as such.

Source: authors' own calculations.

By overlaying and visually comparing detections of images, it appears that some years over-detect and other years under-detect depending on (at least) three factors. First, if an image is of lower quality or has a low contrast, it becomes more difficult for the CNN structure detector to disentangle structures and non-structures, leading to under-detection or over-detection. Second, the time of year of the image influences the saturation of vegetation as well as the colour of the soil. In some instances, trees or bushes can be mistaken for buildings. An attempted quick fix of adding around 1,000 examples of trees to the non-house category when training the algorithm was not sufficient to improve the algorithm's ability to distinguish between these. Third, the time of the day when an image is captured may also influence the number of detections the CNN structure generates. If an image is taken in the middle of the day when the sun is 90 degrees above the ground, shadows are generally not visible, but, if the image is taken in the morning or the afternoon, shadows may be even more noticeable than the differences between e.g. a building and a road. When shadows are prevalent, the general outcome is more detections.

When looking at construction density, the overall trend is broken by the years 2012 and 2014 (Figure 10). For 2012, there is a stark over-detection, which is also visible from manual inspection

of the detections (mainly arising from shadows), whereas, for 2014, many structures are not detected due to the poor quality of the images.

The trend in share of high-quality roofs across the five years, is close to linear, which is largely in line with expectations. This indicates that, although the number of structures detected is sensitive to variation in the quality of the images, the classification of detections into different categories seems to work better.

Although our CNN algorithm has been trained using data from five different cities, and thus five different hours and (possibly) months, it is still not at a stage of development where comparisons over time are trustworthy in all cases. A possible solution to the problem requires, as a minimum, much more training data with much more variation in resolution, contrast, vegetation, and shadows. As noted in section 2.1, most of the training data comes from the city of Tete where all the buildings in the city were tagged, thus allowing for construction of a good dataset of non-structures as well. However, the non-structure category, in particular, clearly needs to be expanded with tags from more cities and many more image types (month, hour, quality) for the comparison of the number of buildings over time to be exact. This will be a topic for future research.

5 Conclusion

This study documented changes to the quantity and quality of housing in an area of around 12 km x 25 km in the northern outskirts of Greater Maputo, from before to after the completion of a major ring road construction project. Two CNNs were adopted to detect and classify buildings on satellite images for the years 2010 to 2018, accessed via Google Earth. The results show that between 2010 and 2018, overall construction density increased from 13 per cent to 50 per cent; the share of high-quality roofs among detections increased from 9 per cent to 24 per cent, and the share of buildings under construction declined from 33 per cent to 26 per cent. Though areas closer to the ring road had higher growth in structures, a higher share of high-quality structures, and a declining share of structures under construction, we cannot convincingly argue that this was driven by proximity to the road rather than just the road being located in a part of town that was growing rapidly. We come to this conclusion as the impact of the road is very different north and south of the city, while there is a consistent pattern of the city growing from the south toward the north.

An important contribution of this study is the realization that the current CNN detections seem sensitive to the input of images, which has ramifications as images usually vary across time in various aspects. Examples of variations include the image quality, time of year, and time of day of image capture. These factors influence the number of detections generated by the algorithm to the point where accurate level comparisons over time become meaningless. Therefore, more training and evaluation should be undertaken. The CNN which classifies buildings into categories may be more reliable than the detector.

References

- Andersen, J.E., P. Jenkins, and M. Nielsen (2015). 'Who Plans the African City? A Case Study of Maputo: Part 1—the Structural Context'. *International Development Planning Review*, 37(3): 329–50

- Avci, O., O. Abdeljaber, S. Kiranyaz, and D. Inman. (2017). 'Structural Damage Detection in Real Time: Implementation of 1D Convolutional Neural Networks for SHM Applications'. In C. Niezrecki (ed.), *Structural Health Monitoring & Damage Detection*, Volume 7. Cham, Switzerland: Springer International Publishing.
- Beam, C. (2019). 'Soon, Satellites Will Be Able to Watch You Everywhere All the Time. Can Privacy Survive?'. MIT Technology Review. Available at: <https://www.technologyreview.com/s/613748/satellites-threaten-privacy/> (accessed 20 December 2019).
- Bernard, L., J. Bird, and A.J. Venables (2016). *The Urban Land Market: A Computable Equilibrium Model Applied to Kampala*. Oxford: University of Oxford.
- Fisker, P., D. Malmgren-Hansen, and T.P. Sohnesen (2019a). 'Using Satellite Data to Guide Urban Poverty Reduction'. Unpublished paper.
- Fisker P., D. Malmgren-Hansen, and T.P. Sohnesen TP (2019b). 'TASIM: Tagged Structures in Mozambique'. (Mendeley Data).
- Fujita, A., K. Sakurada, T. Imaizumi, R. Ito, S. Hikosaka, and R. Nakamura (2017). 'Damage Detection from Aerial Images Via Convolutional Neural Networks'. In Fifteenth IAPR International Conference on Machine Vision Applications (MVA), pp. 5–8. doi: 10.23919/MVA.2017.7986759.
- Government of Mozambique (GoM) (2012). Município lança projecto 'Estrada Circular de Maputo'. Available at: <https://web.archive.org/web/20141018135113/http://www.portaldogoverno.gov.mz/noticias/infraestrur/marco-2012/municipio-lanca-projecto-estrada-circular-de-maputo/> (accessed 20 December 2019).
- Lall, S.V., J.V. Henderson, and A.J. Venables (2017). *Africa's Cities: Opening Doors to the World*. Washington, DC: World Bank.
- Liu, P., L. Xiaoping, L. Mengxi, S. Qian, Y. Jinxing, X. Xiaocong, and Z. Yuanying (2019). 'Building Footprint Extraction from High-Resolution Images via Spatial Residual Inception Convolutional Neural Network'. Remote Sensing. doi: 10.3390/rs11070830.
- Malmgren-Hansen, D., A.A. Nielsen, M.B. Kreiner, R. Saldo, H. Skriver, L.T. Pedersen, J. Lavelle, J. Buus-Hinkler (2019) *High-Resolution Sea Ice Maps with Convolutional Neural Networks*. Available at: <https://orbit.dtu.dk/en/publications/high-resolution-sea-ice-maps-with-convolutional-neural-networks> (accessed 20 December 2019).
- Tiecke, T.G., X. Liu, A. Zhang, A. Gros, N. Li, G. Yetman, K. Talip, S. Murray, B. Blankespoor, and E. B. Prydz. *Mapping the World Population One Building at a Time*. arXiv. 2017. *arXiv preprint arXiv:1712.05839*.
- Wikipedia (2019). 'Estrada Circular de Maputo'. Original by NordNordWest. Available at: https://pt.wikipedia.org/wiki/Estrada_Circular_de_Maputo (accessed 20 December 2019). Licensed under CC BY-SA 3.0 de. Modified by the author.