

Shadow Based Building Extraction from Single Satellite Image

Gurshamnjoy Singh ^{a,b}, Mark Jouppi ^{a,b}, Zhuoran Zhang ^{a,b} and Avideh Zakhor ^{a,b}

^a Signetron Inc. ^b University of California, Berkeley
Berkeley, CA. USA

ABSTRACT

Automatic building extraction in satellite imagery is an important problem. Existing approaches typically involve stereo processing two or more satellite views of the same region. In this paper, we use shadow analysis coupled with line segment detection and texture segmentation to construct rectangular building approximations from a single satellite image. In addition, we extract building heights to construct a rectilinear height profile for a single region. We characterize the performance of the system in rural and urban regions of Jordan, Philippines, and Australia and demonstrate a detection rate of 76.2 - 86.1 % and a false alarm rate of 26.5 - 40.1 %.

Keywords: Building detection, Satellite imagery, Image processing, Remote sensing

1. INTRODUCTION

With the ever-increasing quantity and quality of high resolution satellite imagery, automated methods of analysis and extraction of objects of interest have become increasingly important. Building footprint and height information is important for many applications including image based geolocalization and urban planning. In some urban areas, building footprint information has been manually annotated through crowdsourcing, such as that available through OpenStreetMap. In other places, building dimensions may be available through local government records. However, in many parts of the world, these rich sources of data are not readily available.

One modality to detect buildings is high resolution LIDAR data which can be used to compute digital terrain models (DTM), digital surface models (DSM), and 3D building models.¹⁻⁴ 3D Building models have also been extracted using stereo or multiple images of the same area with aerial and satellite imagery. From the stereo computation, a height can then be estimated for each pixel in the image, and the corresponding height map can be used further for building extraction.^{5,6} For high resolution satellite imagery, building pixels can be classified using morphological image processing techniques.⁷ A morphological building index (MBI) image indicates each pixel's likelihood of being a building. The technique is designed to classify homogeneous and bright groups of pixels falling within an adjustable size range as buildings. This method works particularly well for buildings with bright and uniform roofs, but can miss darker or heterogeneous roofs and assign them lower MBI values. Furthermore, bright parking spots or patches of soil may be incorrectly assigned high MBI values.

A number of building detection approaches use edge detection.⁸ In this paper, we propose a method for building extraction from a single satellite imagery that leverages an edge detection based line segment extraction algorithm. We extract shadow lines corresponding to buildings from a satellite image and couple the shadow line data with texture segmentation and MBI results to construct rectangular approximations to buildings. A ground truth comparison of this method in suburban and urban areas of Jordan is described along with similar analyses in Australia and Philippines. The outline of the paper is as follows: Section 2 details shadow line extraction

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Air Force Research Laboratory, contract FA8650-12-C-7211. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, AFRL, or the U.S. Government.

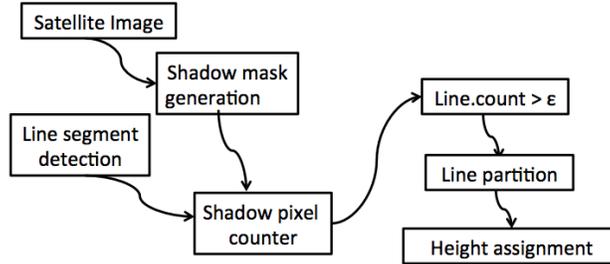


Figure 1: Overall pipeline for shadow building line generation.

from satellite imagery, Section 3 describes shadow line matching with segmentation results and constructing rectangular approximations, Section 4 details pruning for and selecting reasonable building completions, and Section 5 presents experimental results on Jordan and Australia.

2. SHADOW BUILDING LINE EXTRACTION

This section describes the approach used to identify height-attributed building facade lines. Our method detects line segments within satellite imagery and prunes the detected segments using shadow and height attributes. The pipeline used in the process is depicted in Fig. 1, which demonstrates the process by which shadow analysis is used in conjunction with line segment detection (LSD) to detect “shadow building” line segments, a term which refers to lines at the boundary of a building and its shadow. The pipeline is detailed further in the following sections.

2.1 Line segment detection

We use a gradient-based LSD tool to identify all object boundaries characterized by a steep intensity gradient.⁹ As shown in Fig. 2, LSD successfully detects a multitude of object boundaries, a subset of which includes shadow building boundaries. In order to clearly delineate building facades and extract building height information, we isolate line segments bordering building shadows. As detailed shortly, this process uses shadow analysis to determine whether the facade identified by a line segment generates a shadow.

2.2 Shadow mask generation

To segment building shadows, we extract the local histogram peak corresponding to the lowest intensity pixels in the image, a process known as histogram peak clustering.¹⁰ However, satellite imagery typically has a similar intensity profile for vegetation and shadows. Thus, an additional step is needed to eliminate vegetation pixels using the Normalized Difference Vegetation Index (NDVI) prior to histogram peak clustering. Specifically, the satellite imagery is pre-processed in order to remove all pixels for which NDVI exceeds a pre-selected threshold:



Figure 2: (a) Sample building, (b) LSD output superimposed.

$$NDVI = \frac{NIR - VIS}{NIR + VIS} \quad (1)$$

Thresholding the resultant histogram by intensity yields a binary image, which we refer to as the shadow mask. Figs. 3a and 3b display a sample satellite tile and the corresponding shadow mask following pre-processing. Next, using solar azimuth metadata, we identify all line segments preceding shadows while traversing in the direction away from the sun. As displayed in Fig. 1, this process is carried out by the *Shadow pixel counter*, which counts the average number of shadow pixels following the line segment in the direction away from the sun. The *Shadow pixel counter* uniformly samples the line segment and associates a shadow pixel count to each sampled point. The arithmetic mean of all shadow pixel counts along the line segment is referred to as *line.count*. Fig. 4 demonstrates the process pictorially using 3 sample points. Thresholding *line.count* by $\epsilon > 0$ yields a set we term non-partitioned shadow building lines.

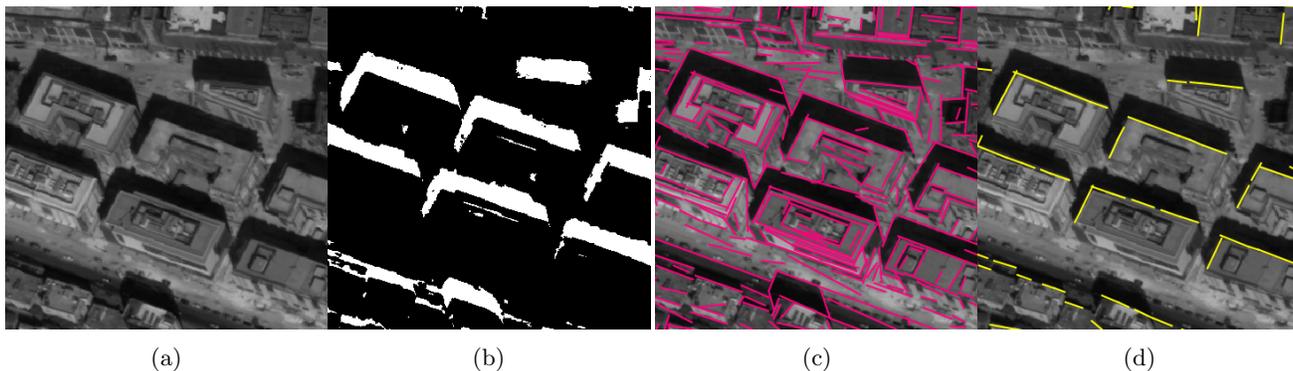


Figure 3: (a) Sample tile. (b) Shadow mask. (c) LSD output. (d) Shadow-building lines.

2.3 Line partition

Given the dense nature of buildings in urban settings, it is likely for detected line segments to border multiple buildings. This complicates the problem at hand as an ideal one-to-one mapping between shadow building line segments and buildings is required for accurate building completion.

Fig. 5a displays shadow building line segments superimposed on a region with high building density. The isolated line segment overlaps multiple distinct buildings in the region. Notably, the building boundary is marked by a sharp transition in shadow length. Thus, we propose a *line partition* algorithm to detect distinct building boundaries by thresholding the first derivative of the shadow length profile along the line segment as shown in Figs. 5b and 5c. Fig. 5d displays the computed point of partition. The line segment is then partitioned into two different line segments at the point of partition.

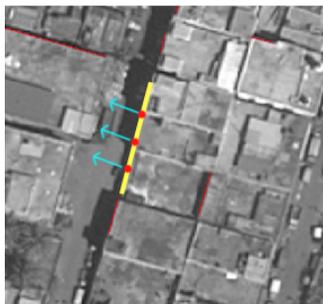


Figure 4: Methodology for shadow pixel counting. Yellow line: line segment in question. Red points: points sampled on the line segment. Cyan arrows: Direction in which shadow pixels are counted. The final shadow pixel count is the mean count from all sampled points.

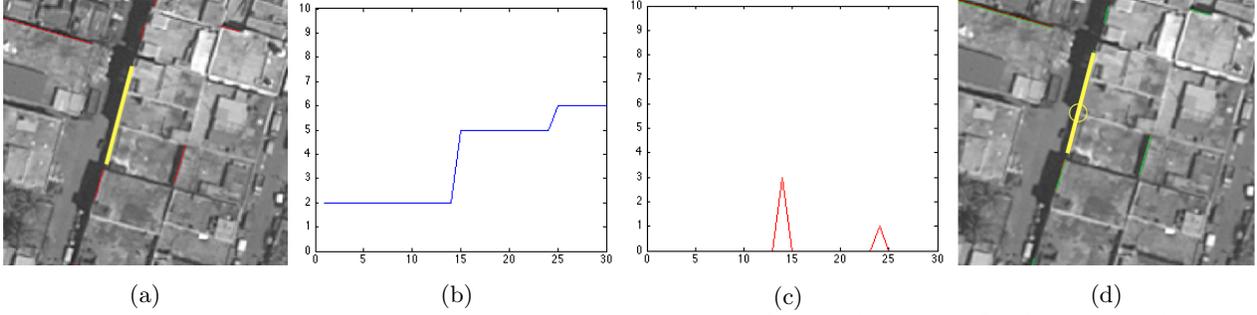


Figure 5: (a) Line segmented bordering multiple buildings. (b) Sample height profile for the highlighted line segment. (c) First derivative of the height profile. (d) Point of partition for the line segment. Note: The height-profile data is self-generated for clarity.

This technique generates shadow building lines with a reasonable guarantee of a one-to-one mapping to buildings. Additionally, Figs. 3c and 3d display the LSD result and shadow building lines for the sample tile shown in 3a.

2.4 Height assignment

We assign a shadow length value to each shadow building facade line. Using sun elevation data as shown in Fig. 6, we are able to compute a region height in the vicinity of the shadow building line segment:¹¹

$$h = \frac{l}{\tan(\phi)} \quad (2)$$

where h is the region height, l is the shadow length, and ϕ is the sun zenith angle. This procedure allows us to construct an approximate height map using purely shadow information. Appended with completed rectangular approximations to buildings, our technique generates a 3D-height map reconstruction of the region using only 1 satellite image, making our procedure suitable for modeling height profiles of regions with low satellite coverage.

3. SHADOW BUILDING LINE RECTANGULAR COMPLETION

We now describe our method of rectangular building approximation leveraging the aforementioned techniques. Having detected shadow building line segments, the next step is to determine the extent of the building corresponding to that line. We apply texture segmentation to decompose the image into approximate logical regions and then analyze each region. Specifically, among all regions touching a given shadow building line, we choose and assign one region that is most likely to be a building. In this way, for each shadow building line, a region with the greatest likelihood of being the corresponding building is found. Rectangular approximations of buildings can then be computed from the segments.

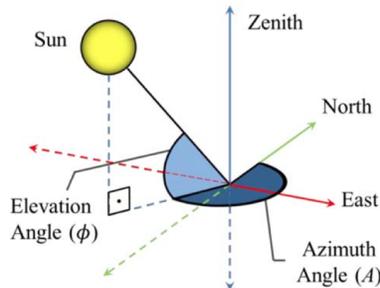


Figure 6: Solar angle notation.¹¹

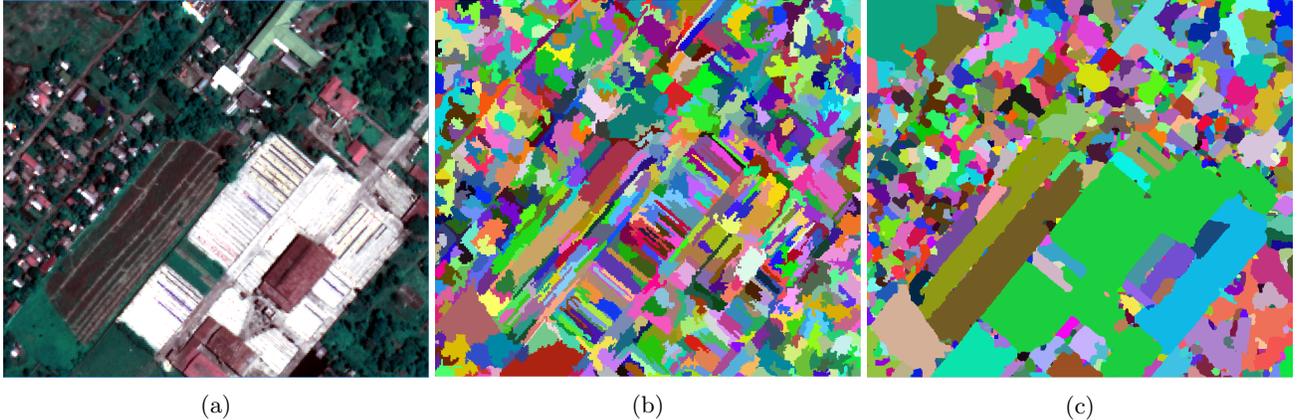


Figure 7: (a) Original Philippines tile. (b) Mean shift clustering. (c) Multiscale combinatorial grouping.

Our method makes the inherent assumption that buildings are of a given size range, as determined by the general size of the regions obtained via texture segmentation. This could be remedied to some extent by allowing the selection of multiple contiguous regions falling within matching criteria to match a given shadow building line.

3.1 Segmentation

The first step is to perform texture segmentation on the satellite view. The goal is to break the image up into distinct logical regions such that each building is a separate object. Then for each shadow building line we can determine the nearby region most likely to correspond to a building. We employ two texture segmentation techniques and compare them in Fig. 7. For mean shift clustering (MSC), we control the *spectralRadius*, *spatialRadius*, and *minimumSegmentSize*. Specifically, *spectralRadius* defines the spectral interval; larger values require less spectral separation between regions. Thus, a larger *spectralRadius* produces larger regions. *SpatialRadius* indicates the neighborhood to search within when creating regions. A larger *spatialRadius* thus produces larger regions. *MinimumSegmentSize* denotes the smallest segments allowed such that segments smaller than this are forced to merge with others until they are at least this large.

For multiscale combinatorial grouping (MCG), the only parameter we control is the threshold k for the Ultrametric Contour Map used to compute the boundaries of segmentation at any particular scale. Larger k values result in generally larger segments, but there is no minimum segment size constraint as in MSC. MCG segmentation is smoother as shown in Fig. 7, and also much faster than MSC, but there is no easy way to translate k values between different satellite images, such that a reasonable k value for one image may yield a segmentation of greatly different quality on another image.

Results of texture segmentation optimized for buildings in our Jordan satellite image tile in Fig. 8a are shown in Fig. 8b, colored to show the spectral mean for each segment. This segmentation was produced using MSC with the following parameters: *spectralRadius* = 15, *spatialRadius* = 15, *minimumSegmentSize* = 200. For our Philippines tile in Fig. 7a, we show MCG segmentation results tuned for our purposes in Fig. 7c, generated with $k = 0.13$.

3.2 Region evaluation

Each region is evaluated based on its properties to determine its likelihood of being a building. This is essentially object-based classification, but rather than classifying each object independently, we assign each object a likelihood for being a building. We experimented with various metrics to determine this including geometrical properties such as length width ratio, perimeter area ratio, and rectangular fit in addition to MBI and NDVI.

We have empirically found MBI to be the best predictor for buildings among the various properties, so we simply use MBI as a descriptor for likelihood of an object being a building. We use shape and NDVI properties to filter out regions in a later stage of the pipeline to be described in Section 4.1. Fig. 8c and 8d show MBI and NDVI for the segmentation shown in Fig. 8d.

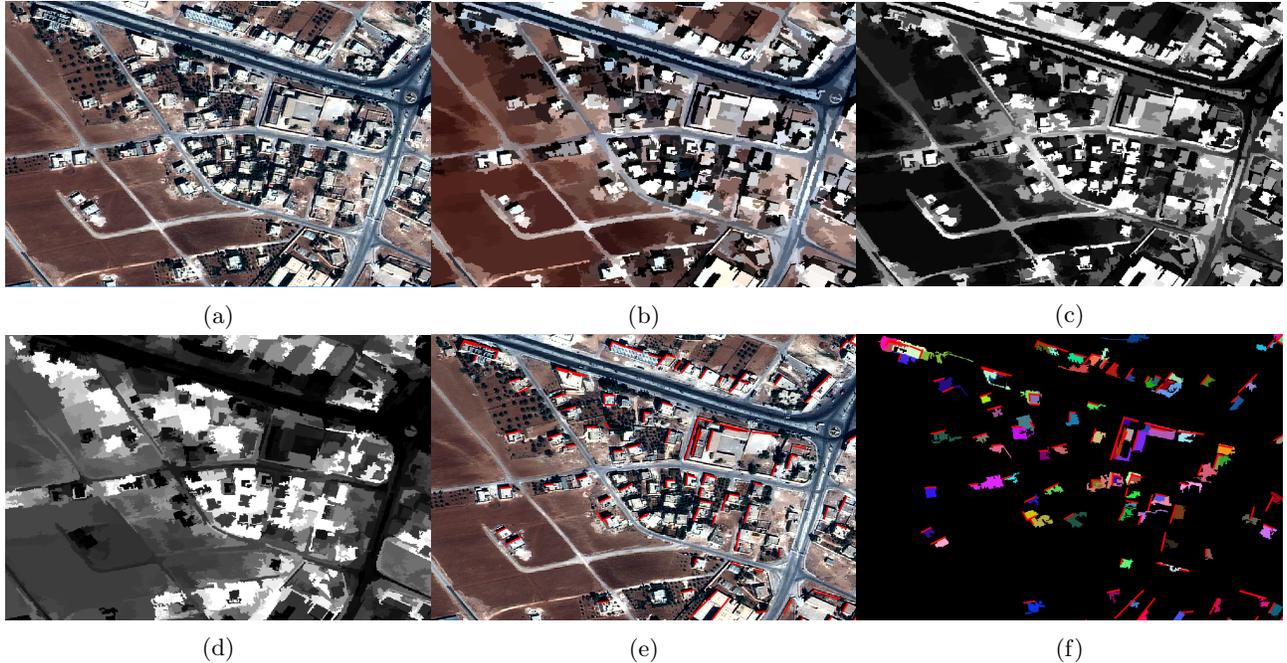


Figure 8: (a) Original tile. (b) Segmentation. (c) MBI per segment. (d) NDVI per segment. (e) Shadow building lines on original image. (f) Selected regions for each shadow building line.

3.3 Shadow building line-segment intersections

To determine the candidate texture segments for each shadow building line, the intersections between shadow building lines and segments are computed. Once each segment has been evaluated, and intersecting segments for each shadow building line have been determined, we choose the most likely building segment for each shadow building line. Fig. 8f shows examples of the final chosen segments for each shadow building line compared with the original image. As seen, performance is reasonably good with only a few missed segments for shadow building lines. Also, few cases of undersegmentation result in smaller segment matches.

Following the matching process, a rectilinear approximation is constructed for each matched building segment, which we refer to as a “building completion. An example of this is shown in Fig. 8. This process selects a shadow building line segment and computes the maximum building segment length in the perpendicular direction. The algorithm samples the line segment at various points and selects the maximum building segment length following the line defined by the point and the slope of the perpendicular line segment. This length, therefore, resolves

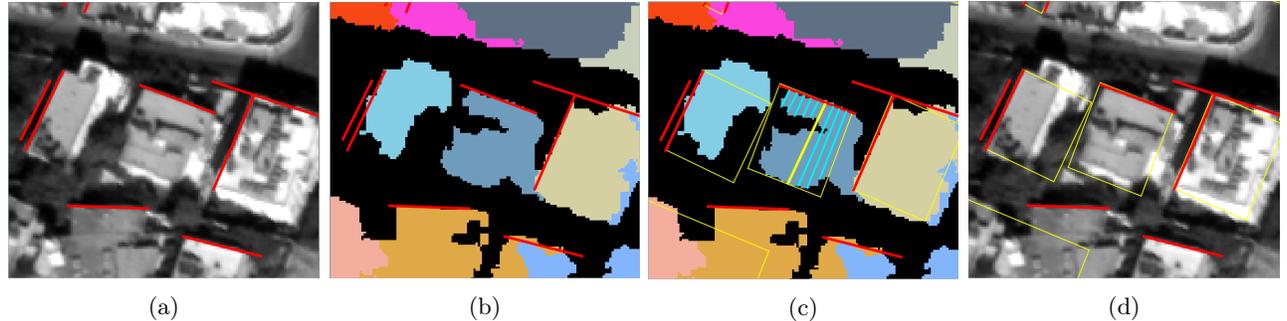


Figure 9: (a) Shadow building lines on original tile. (b) Shadow building lines on segment tile. (c) Perpendicular lines (cyan) extended from sampled points on the shadow building line (maximum length extension along with building completion identified in yellow). (d) Building completion on original tile. Note: the perpendicular extensions are robust against small gaps in the segmentation results.

the missing dimension necessary for building completion.

4. BUILDING COMPLETION PRUNING

Due to over-detection by LSD and shadow building line extraction, building completion results are subject to a high false positive rate, corresponding to building completions in areas of no buildings, and redundancy, corresponding to multiple completions for one given building. Thus, we employ a pipeline that prunes building completion to retain the most reliable ones. Fig. 10a displays the corresponding pipeline, which uses a global pruning step in order to disqualify completions that do not satisfy hard thresholds and a graphical pruning step which selects the most reliable completions among many.

Fig. 11a displays a region from urban Jordan in which building completion results display both of these attributes. As observed, building completions exhibit variable density depending on the region of the tile. Next, we develop heuristics to prune redundant and false completions.

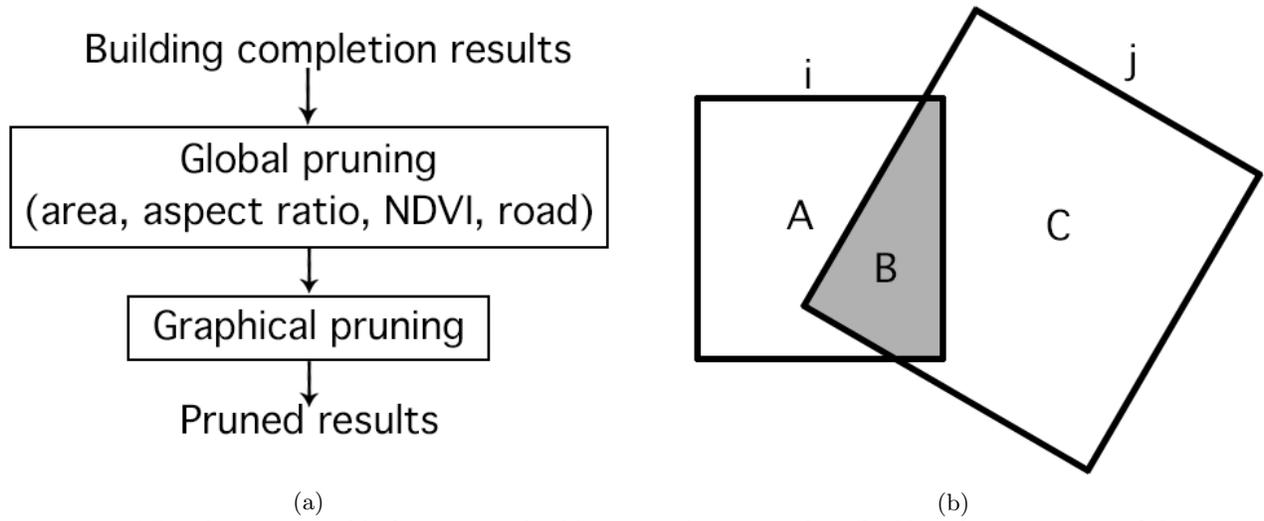


Figure 10: (a) Pipeline responsible for pruning building completion results. (b) Visual representation of the scoring function.

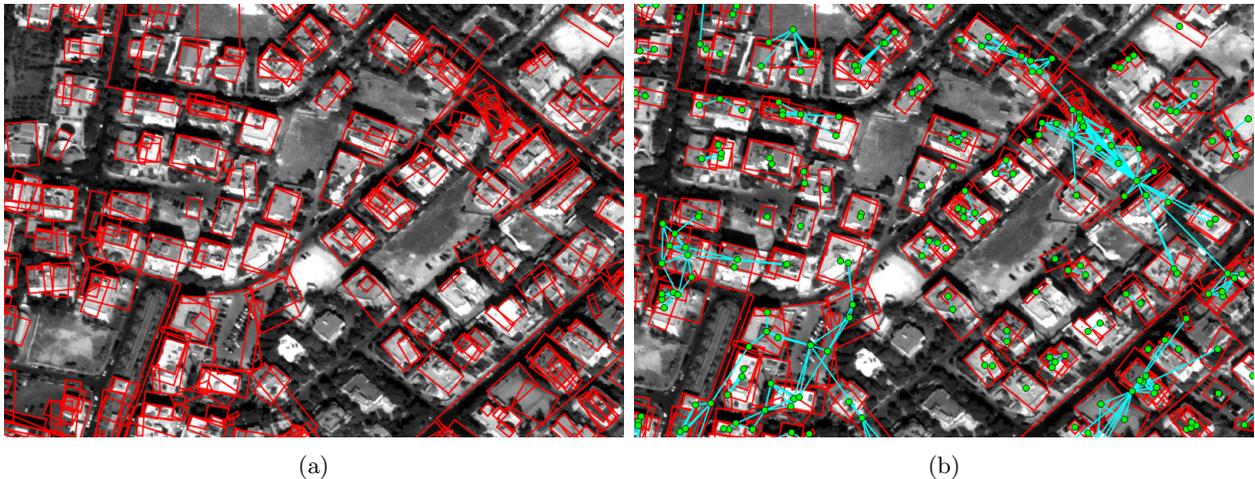


Figure 11: (a) Entire set of building completions for an urban region. (b) Graphical display of urban region. Green: nodes, Cyan: edges.

4.1 Pruning overlapping building completions

Prior to implementing local and region-dependent pruning techniques, we make use of global filtering techniques to eliminate non-realizable building completions. Specifically, we implement area filtering under the assumption that building completions in the given region do not deviate significantly in area from the population mean. We prune all building completions deviating from the populating mean, μ , by two standard deviations, 2σ . Under the assumption that ground-truth building structures have dimensions of approximately same magnitude, we prune all building completions with aspect ratio, given by $\min\{\frac{l}{w}, \frac{w}{l}\}$ where w = width and l = length, less than the threshold, ϵ_{aspect} . By treating building rooftops and vegetation as separate classes, we also prune completion results that cover vegetation. Strictly speaking, if the percentage of pixels in the completion result exceeds a vegetation threshold, ϵ_{NDVI} , the building completion is pruned.

Open Street Map (OSM) also proves to be a ready source of data to prune false alarms in many world regions. If any road line in OSM crosses a building completion, we consider the two dissected completion polygons and compute the ratio between the area of the larger polygon to the smaller polygon. Thresholding this ratio is an effective means of determining and pruning completions that lie on the center of a road.

4.2 Graphical Based Pruning

We construct a graphical representation in order to remove redundant overlapping building completions. The graph is denoted by $G = (U, E)$, where U is the set of nodes corresponding to building completions, and E is the set of edges corresponding to building overlap. Let, $u_i \in U$ refer to building completion i . The edge (u_i, u_j) is a directed edge from node u_i to node u_j with a weight defined by the scoring function, $s(i, j)$, denoting the overlap between i th and j th building completion, defined as follows:

$$s(i, j) = \frac{|b_i \cap b_j|}{|b_i|} \quad (3)$$

where b_i is the set of pixels enclosed by building completion i .

Fig. 10b shows a visual example of the scoring function. In this case:

$$s(i, j) = \frac{B}{A + B}$$

$$s(j, i) = \frac{B}{B + C}$$

Note that the s is not necessarily commutative, as is the case with building completions of different areas. Fig. 11b displays the graphical representation of the urban region in Fig. 11a. For ease of visualization, directed edges (u_i, u_j) and (u_j, u_i) are reduced to one undirected edge $\{u_i, u_j\}$. As noted in the display, the graph contains regions of high connectivity along with regions of low connectivity. Here, we associate the phrase ‘‘connectivity’’ to the density of edges. Pruning building completions, thus, corresponds to pruning nodes to reduce the overall connectivity of the graphical representation.

We proceed by adopting a connected-component based algorithm that prunes nodes based on local connectivity of the graphical structure. In order to avoid impacting detection rate, we must distinguish between building completions with high or low overlap. Intuitively, completions with high overlap correspond to multiple detections for the same building object. On the other hand, completions with low overlap correspond to independent building detections. To quantitatively define regions of high and low overlap, we make use of the overlap scoring function, s . Precisely, building completion i and j have high overlap iff $\max\{s(i, j), s(j, i)\} > \epsilon$ and low overlap otherwise. Fig. 12a displays the overlap regions with respect to the scoring coordinates.

In the case of high overlap between two building completions, we eliminate one. For pruning purposes, each building completion, i , is assigned a score, l , using the following metric:

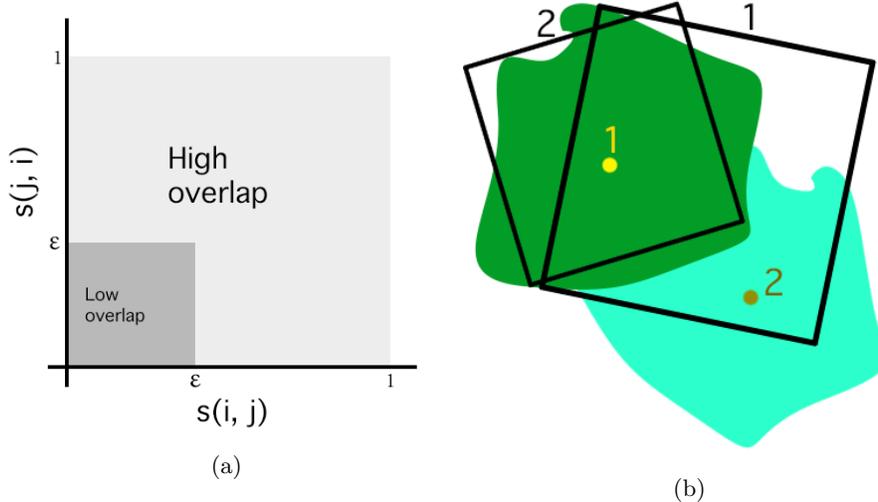


Figure 12: (a) Quantitative regions corresponding to building overlap. (b) Visual example for scoring scheme.

$$l(i) = \sqrt[n]{\sum_{j \in O} p_j^n} \quad (4)$$

where the set O indexes all building segments that underlay the i^{th} building completion and p_j refers to the percentage of pixels occupied by the j^{th} building segment in building completion i . Fig. 12b displays a simple example in order to demonstrate the scoring scheme. For completion 1, $O = \{1, 2\}$, $p_1 = 0.5$, $p_2 = 0.3$. Thus, for $n = 2$, we obtain $l(1) = \sqrt{(0.5)^2 + (0.3)^2} = 0.58$. For completion 2, we obtain $O = \{1\}$, $p_1 = 0.9$. For $n = 2$, we have $l(1) = \sqrt{(0.9)^2} = 0.9$.

The scoring scheme in Equation 3 is selected in order to minimize a large number of distinct overlaid building segments by a building completion, $|O|$, and maximize the percent overlap between an overlaid segment and the given completion. Therefore, this scheme discourages completions that could overlap multiple buildings and encourages completions that well overlap and enclose a building segment. Empirical measurement trials on scoring data show that detection rate and false positive rate are relatively independent of n for $n \geq 2$ in Equation 3. Therefore, we select $n = 2$ for all scoring metric computations.

Given the highly dense nature of building completions, we devise an algorithm to prune excessive completions and preserve the most logical results. To begin, we examine an entangled set of overlapping building completions and proceed by selecting the best fit representation from two completions that overlap the most. The algorithm proceeds in an iterative manner until no set of building completions with a high overlap remain.

To identify regions to prune, our algorithm performs a connected component analysis on the graphical representation. Each connected component is treated independently to only prune high overlap completion results.

The pseudo-code is as follows:

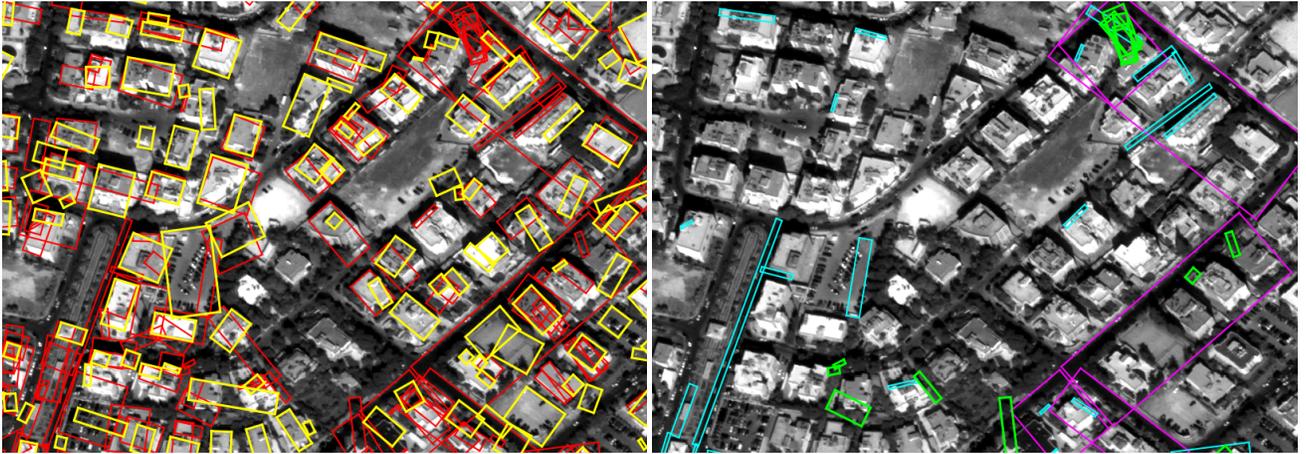
```

 $U \leftarrow$  Nodes with area, aspect ratio, ndvi pruning
 $E \leftarrow$  Edges s.t.  $\max\{s(i, j), s(j, i)\} > \epsilon$ 
 $CC \leftarrow$  Index of connected components
for  $cc$  in  $CC$ :
     $E(cc) \leftarrow$  edges sorted by weight
    for  $e = \{i, j\}$  in  $E(cc)$ :
        completion  $\leftarrow \arg \min\{l(i), l(j)\}$ 
         $U \leftarrow$  Remove completion
         $E \leftarrow$  Remove completion edges
output  $U$ 

```

Iteratively, the algorithm selects a set of building completions with a high overlap, identified as a connected component, and proceeds by selecting two completions with the highest overlap. Next, the completion with the higher score, l , is selected with the other completion removed and the graphical representation is updated. Following each iteration, we see that $|E|$, the number of edges, decreases by at least one. Thus, our algorithm constructs the final completion output after at most $|E|$ iterations.

Fig. 13a displays the output of the post-processed building completion results overlaid on the entire set of building completion results. Notably, filtering and graphical pruning eliminates a large set of false positive building completion results.



(a)

(b)

Figure 13: (a). Yellow: Filtered and pruned completions. Red: Total set of completions. (b). Cyan: Aspect ratio pruning. Magenta: Area pruning. Green: NDVI pruning.

5. RESULTS

The building approximation pipeline was tested on urban and rural settings in Jordan. The parameters used for characterization are:

- *Detection rate*: Ratio between number of buildings detected by the method and the ground truth.
- *False positives*: Ratio between number of falsely reported buildings by the method and total number of buildings reported by the method.

Building detection in the two urban regions was measured using the intersection over union (IOU) metric with a threshold of 50 %. The metric accounts for a detection if the intersection of the detection and ground truth over the union of the two is larger than the threshold.

In the rural region, ground truth was in the form of points representing the centroids of buildings. Building completions were then matched to ground truth points 1-to-1 as follows: if a ground truth point is inside the bounding box given by a completion, the two are paired. A second pass is made to pair ground truth points and completions where the centroid of the completion and the ground truth point are within a predetermined tolerance of each other. At the end of this process, completions that are not paired are considered false alarm, and ground truth points that are not paired are considered missed.

Table. 1 displays the accuracy results for the entire pipeline generated for two urban regions, named A and B, in Jordan as depicted in Figs. 14a and 14c, one rural area in Jordan named C as depicted in Fig. 14e, and one rural area in Australia as depicted in Fig. 14g. The pruned building completions for the tiles in Figs. 14a, 14c, 14e, and 14g are shown in Figs. 14b, 14d, 14f, and 14h respectively. From Table 1 and Fig. 14, we conclude that the false alarm rate is comparatively higher in urban regions than rural regions in Jordan. This can be attributed to the high density of buildings in the rural area whereby shadows of buildings on one side of the street fall on top of the buildings on the other side; another reason might be excessive overlap of rectilinear building approximation where the graph based pruning of Section 4.2 might not manage to untangle and prune them in an optimal way.

Another observation from Table 1 is that the detection rate and false positive rate performance for Australia is considerably worse than in Jordan due to excessive vegetation. This is due to the fact that the overall pipeline relies heavily on approximating the shadows of buildings through a shadow mask and segmentation of their shape. In some regions of Australia, tall trees grow over the tops of buildings, partially obscuring the geometry of the rooftop as seen in Fig. 16b. Since the shape of the top of the tree is rarely regular, the visible portion of the rooftop also becomes irregular, resulting in an erroneous rectilinear approximation. It is possible that enough rooftop pixels are covered by vegetation that the remaining rooftop pixels fail to satisfy the aspect ratio or perimeter-to-area constraints in building completion. Another example of vegetation interfering with detection rate is shown for rural area of Philippine in Fig. 16a. As seen, the presence of large amounts of vegetation behind buildings reduces the number of dark pixels visible in the satellite image that should be classified as shadow pixels. Since the algorithm relies heavily on detecting a shadow building line for each building, absence of shadow pixels where a precludes many buildings from being detected.

Table 1: Statistical results for regions in Jordan and Australia

Region	Type	Area (sq. km)	Detection	False Positive	Num. Buildings
Jordan (A)	Urban	0.5	86.1 %	40.1 %	374
Jordan (B)	Urban	2.2	62.5 %	38.0 %	1288
Jordan (C)	Rural	4.9	76.2 %	26.5 %	628
Australia	Rural	3.3	46 %	67 %	30



(a)

(b)



(c)



(d)



(e)



(f)



(g)



(h)

Figure 14: (a) Urban area A in Jordan. (b) Pruned completions for A. (c) Urban area B in Jordan. (d) Pruned completions for B. (e) Rural area C in Jordan. (f) Rural area in Australia. (g) Pruned completions for image in part f.

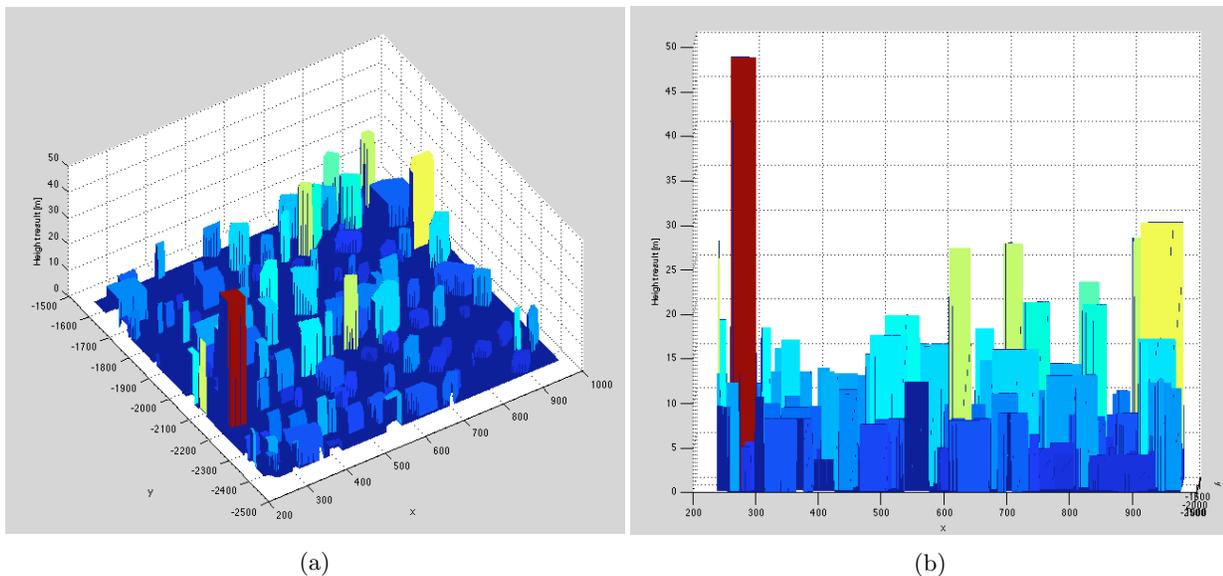


Figure 15: (a) Extracted three dimensional height profile of the region. (b) Side view of the height profile of the region.



Figure 16: (a) Philippines Tile. (b) Australia Tile.

Given constructed building approximations and the corresponding height attributes, we also construct a three-dimensional rendering of the satellite tile. A view of the rendering, corresponding to the region in Fig. 14a is displayed in Fig. 15. The rendering captures a compact representation of the geographical scene with building shapes and heights. Such reconstructions can be employed in 3D scene matching algorithms in which a similar representation can be extracted from ground level query images.

The above techniques are parallelizable since they are identically applicable to small tiles of a satellite image rather than the entire image. We use GDAL utility `gdal.translate` to tile a satellite image into a set number of tiles, generally the number of cores we wish to process on, and simply execute a copy of the pipeline on each tile for low-hanging gains in speed.

Fig. 17 shows a 9.9 sq. km satellite tile from Australia. The image dimensions are 20073 x 9490 and the parallel pipeline completed in 7.5 hours on 6 core with 25 megapixel subtiles, with slight overlap to make sure

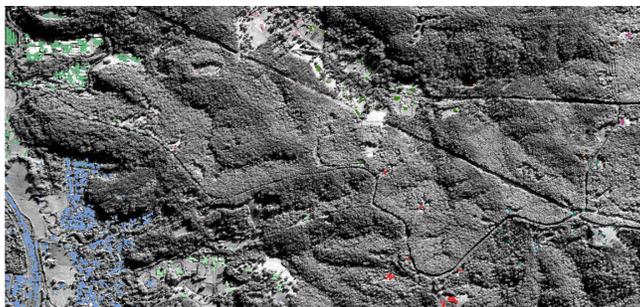


Figure 17: Large region from Australia, with building detections color coded by the subtitle they came from in parallel processing.

no buildings are cut by the boundaries. A process such as the graph pruning described above may be used to resolve the boundary effects, though it was not implemented for the results depicted in Fig. 17. No ground truth currently exists for this tile to obtain performance statistics.

6. CONCLUSION

In this paper, we outlined automated methods for building detection in a single satellite image. Our method’s pipeline uses shadows as a primary mean for measuring building heights and rectilinearly approximates building objects via texture segmentation and line segment detection.

The developed techniques rely solely on a single satellite image to execute the entire pipeline. This offers a great advantage over traditional building detection methods that require rich remote sensing data such as LIDAR or multiple overlapping satellite imagery. Given the scarce nature of satellite imagery and other remote sensing data in many parts of the world, our method offers a starting point for various geolocalization algorithms that would be otherwise deficient in input.

REFERENCES

- [1] Rottensteiner, F., Trinder, J., Clode, S., and Kubik, K., “Building detection using lidar data and multi-spectral images,” in [*Digital Image Computing: Techniques and Applications*], **2**, 673–682, CSIRO (2003).
- [2] Rottensteiner, F. and Briese, C., “A new method for building extraction in urban areas from high-resolution lidar data,” *International Archives of Photogrammetry Remote Sensing and Spatial Information Sciences* **34**(3/A), 295–301 (2002).
- [3] Verma, V., Kumar, R., and Hsu, S., “3d building detection and modeling from aerial lidar data,” in [*Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*], **2**, 2213–2220, IEEE (2006).
- [4] Sohn, G. and Dowman, I., “Data fusion of high-resolution satellite imagery and lidar data for automatic building extraction,” *ISPRS Journal of Photogrammetry and Remote Sensing* **62**(1), 43–63 (2007).
- [5] Paparoditis, N., Cord, M., Jordan, M., and Cocquerez, J.-P., “Building detection and reconstruction from mid-and high-resolution aerial imagery,” *Computer vision and image understanding* **72**(2), 122–142 (1998).
- [6] San A, D. K. et al., “Automatic building extraction from high resolution stereo satellite images,” (2007).
- [7] Huang, X. and Zhang, L., “Morphological building/shadow index for building extraction from high-resolution imagery over urban areas,” *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of* **5**(1), 161–172 (2012).
- [8] Wei, Y., Zhao, Z., and Song, J., “Urban building extraction from high-resolution satellite panchromatic image using clustering and edge detection,” in [*Geoscience and Remote Sensing Symposium, 2004. IGARSS’04. Proceedings. 2004 IEEE International*], **3**, 2008–2010, IEEE (2004).
- [9] Rafael Grompone von Gioi, Jeremie Jakubowicz, J.-M. M. and Randall, G., “Lsd: A fast line segment detector with a false detection control,” *Transactions on Pattern Analysis and Machine Intelligence* **32**(4), 722–732 (2010).

- [10] Wei, Y., "Urban building extraction from high-resolution satellite panchromatic image using clustering and edge detection," *Geoscience and Remote Sensing Symposium* **3** (2004).
- [11] Ali Ozgun Ok, C. S. and Yuksel, B., "Automated detection of arbitrarily shaped buildings in complex environments from monocular vhr optical satellite imagery," *Geoscience and Remote Sensing Symposium* **51**(3) (2013).