

Article

Earth Observation and Machine Learning to Meet Sustainable Development Goal 8.7: Mapping Sites Associated with Slavery from Space

Giles M. Foody ^{1,2,*}, Feng Ling ³, Doreen S. Boyd ^{1,2}, Xiaodong Li ^{1,3}
and Jessica Wardlaw ^{1,2}

¹ School of Geography, University of Nottingham, Nottingham NG7 2RD, UK; doreen.boyd@nottingham.ac.uk (D.S.B.); lixiaodong@whigg.ac.cn (X.L.); jessica.wardlaw@nottingham.ac.uk (J.W.)

² The Rights Lab, Highfield House, University of Nottingham, Nottingham NG7 2RD, UK

³ Key Laboratory of Monitoring and Estimate for Environment and Disaster of Hubei Province, Institute of Geodesy and Geophysics, Chinese Academy of Sciences, Wuhan 430077, China; lingf@whigg.ac.cn

* Correspondence: giles.foody@nottingham.ac.uk; Tel.: +44-115-9515430

Received: 14 January 2019; Accepted: 24 January 2019; Published: 29 January 2019



Abstract: A large proportion of the workforce in the brick kilns of the Brick Belt of Asia are modern-day slaves. Work to liberate slaves and contribute to UN Sustainable Development Goal 8.7 would benefit from maps showing the location of brick kilns. Previous work has shown that brick kilns can be accurately identified and located visually from fine spatial resolution remote-sensing images. Furthermore, via crowdsourcing, it would be possible to map very large areas. However, concerns over the ability to maintain a motivated crowd to allow accurate mapping over time together with the development of advanced machine learning methods suggest considerable potential for rapid, accurate and repeatable automated mapping of brick kilns. This potential is explored here using fine spatial resolution images of a region of Rajasthan, India. A contemporary deep-learning classifier founded on region-based convolution neural networks (R-CNN), the Faster R-CNN, was trained to classify brick kilns. This approach mapped all of the brick kilns within the study area correctly, with a producer's accuracy of 100%, but at the cost of substantial over-estimation of kiln numbers. Applying a second classifier to the outputs substantially reduced the over-estimation. This second classifier could be visual classification, which, as it focused on a relatively small number of sites, should be feasible to acquire, or an additional automated classifier. The result of applying a CNN classifier to the outputs of the original classification was a map with an overall accuracy of 94.94% with both low omission and commission error that should help direct anti-slavery activity on the ground. These results indicate that contemporary Earth observation resources and machine learning methods may be successfully applied to help address slavery from space.

Keywords: slavery; mapping; crowdsourcing; object-based target detection; convolutional neural network

1. Introduction

It is estimated that over 40 million people in the world can be classed as modern slaves, unable to leave or refuse exploitative activity [1]. Ending slavery has been a goal of numerous agencies for centuries. Anti-slavery activity was further supported in 2016 with the launch of the United Nations Sustainable Development Goals (UN SDGs), specifically goal 8.7, which seeks to promote full productive employment and decent work for all and end modern slavery by 2030 [2,3]. These laudable

activities, however, require accurate information and firm evidence on slavery and this is often unavailable for slavery is “hidden”. For example, basic information on slave numbers and their location is required if direct action to liberate those enslaved or to support policy change and other intervention activity is to be successful. Whilst in many instances, it may be challenging to determine if any individual is enslaved, sometimes the industries in which they work have characteristic properties that can be observed and thus used to focus anti-slavery actions.

Some industries that are associated with slave labour may be observable from space. For example, fish farms established in mangrove forests [4] or excavations associated with mining [5] can be identified in fine spatial resolution satellite sensor images. Indeed, it is likely that a significant proportion of slavery can be observed from space. In Asia, a considerable proportion of the workforce used in the manufacture of bricks is believed to be enslaved [6]. The lack of effective preventative action and prosecution of those abusing the slaves often allows the exploitative practices, commonly involving migrants from economically and socially marginalised communities, to continue [6–8]. The size and location of this slave population are, however, not known. Estimates of slave numbers vary greatly. For example, in India, it has been estimated that approximately 70% of the workforce in brick kilns used to manufacture bricks is forced labour. In the Punjab of Pakistan, it is suggested that ~40% of the workforce is enslaved, working within bonded labour practices [9]. Many of the slaves are children, and some 250,000 children may work within Pakistani brick kilns [10,11]. The effort to meet the UN SDG 8.7 can be aided by accurate and up-to-date information on the number and geographical distribution of brick kilns. Although the Brick Belt is well-known as a region associated with slavery, information on the number and location of brick kilns is poor and uncertain.

As an initial step to providing information to inform anti-slavery efforts in the Brick Belt, fine spatial resolution satellite sensor data have been used to make the first statistically credible estimate of the total number of brick kilns contained in its ~1,550,000 km² area [12]. This estimate, centred at approximately 55,387 brick kilns, suggests a large slave population exists within the region. The results of [12] suggest that brick kilns are not uniformly distributed across the Brick Belt, with kiln density varying greatly across the region although the sample size used limits detailed assessment of the issue. In addition, the work of [12] was deliberately based upon simple methods that could be deployed easily and at low or no cost by bodies with little specialist knowledge. In particular, the brick kilns were identified using a potentially subjective visual interpretation of satellite imagery. Although a step-forward in the journey to meet SDG 8 and a contribution to initiatives such as Alliance 8.7 (www.alliance87.org), this approach can take considerable time and requires substantial human input if large areas are to be studied [13]. Moreover, although this work has been useful in itself, and one brick kiln highlighted by [12] was raided leading to the release of 24 slaves, the work of anti-slavery agencies would particularly benefit from information on the actual locations of each of these brick kilns. Thus, in this paper, the aim is to explore the use of contemporary geoinformation and machine learning methods for the automated mapping of brick kilns, to allow rapid, repeatable and rigorous study of an industry associated with slavery that anti-slavery organisations can use to inform their work. This work forms a step toward the development of an approach that could be used to allow wall-to-wall mapping of the entire Brick Belt on a repeatable basis.

Recent technological developments in three somewhat inter-related areas offer tremendous potential for the mapping of brick kilns over large areas. First, fine spatial resolution satellite sensor data in which brick kilns should be identifiable are now widely available. While much environmental monitoring has made use of relatively moderate spatial resolution (e.g., 30 m) systems such as those carried by the Landsat satellites, the provision of fine, sub-meter, spatial resolution imagery pioneered by IKONOS in 1999 has revolutionised environmental sensing. This is especially the case because such imagery is freely available via resources such as Google Earth, which began providing imagery from 2001. The latter has been used in a wide array of application areas and includes continuously updated remote sensing imagery from various sources mostly from the year 2000 [14]. Critically, as well as using freely available Landsat sensor data, Google Earth also provides high-resolution imagery, notably

from systems such as Digital Globe's WorldView systems, which provide multispectral imagery with sub-meter resolution.

The second key technological development that supports this study is the rise of citizen sensing as a source of ground reference data [15–17]. The full potential of remote sensing as a source of information often hinges on the availability of high-quality reference data to help train image analyses and to evaluate the quality of predictions. Reference data are, however, challenging to acquire. However, the rise of citizen sensing, which takes many forms and may be described by various terms [18], can now provide ground reference data to support analyses of remotely sensed images. Although citizen science has a long history, it is only since the proliferation of mobile devices that are location aware and the development of web 2.0 technology that its use for ground reference data collection in support of remote sensing studies has become practical. This radical change has facilitated the development of the broad area of volunteered geographic information (VGI) that is impacting greatly on geography and geospatial science as a whole [19]. VGI for ground reference data can be collected in various ways. The data collection could be an active process steered by researchers or passive, making use, for example, of information freely available on social media sites. This study followed an active strategy, with citizen sensors or volunteers provided with images from selected locations via an internet platform. The platform used here was the Zooniverse platform [20] launched in 2009. From earlier work, it is known that volunteers can provide accurate annotations of brick kilns [12]. As based on VGI, the process is, however, fraught with challenges ranging from concerns over data quality through to the ability to sustain an active set of high-quality contributors in time [21].

The third and final technological development of relevance to this study is the recent development of advanced image classifiers. The classification methods used in remote sensing have increasingly moved away from conventional statistical classifiers to machine learning methods, with numerous studies indicating the latter can yield more accurate classifications [22–24]. In addition, there have been considerable developments in object- rather than pixel-based image analysis tools [25]. Methods for automated image analysis have progressed greatly in the recent past. For example, the ImageNet challenge 2012 was won with a convolution neural network (CNN) [26]. A CNN is a type of deep feedforward neural network, which has been found to be effective for the analysis of imagery. The basic CNN approach was further developed into region-based, R-CNN, an enhanced classifier but one that is slow [27]. Work to enhance the speed of analysis resulted in the Fast R-CNN and then in 2015 the Faster R-CNN [28]. The Faster R-CNN takes an image as input and generates from it a set of rectangular object proposals that are sub-sets of the image that are predicted to contain within them an object of interest at a stated degree of membership (e.g., the probability of the proposal containing a brick kiln). The potential of these CNN-based classifiers has been recognised in remote sensing, and a range of classifiers have been applied to remote sensing images for the detection of various object types (e.g., [29–34]). The output of a region-based CNN classification is typically a set of proposals that potentially contain a target of interest set within a rectangular bounding box that fits closely around the target together with an estimated measure of the strength of membership to the target class (e.g., a probability of belonging to the target class).

The three main technological resources used in this work are all highly contemporary, and just a decade ago this work would have been impossible. Here, it is hoped they can be used to shape a method for routine and inexpensive monitoring to aid addressing UN SDGs. The aim of this paper is to explore the potential offered by these technological developments to map the locations of sites such as brick kilns, in order to focus anti-slavery resources and activity effectively.


2. Materials and Methods

The research focused on Bull's Trench brick kilns as these are known to be associated with slave labour [6]. These kilns have a characteristic shape, typically oval although some are circular and often with a tall chimney in the middle, see Figure 1. The size of the kilns varies but, as a guide, the radius of circular kilns within the region is often in the order of 33 m. Both of these kiln types exist across the

“Brick-Belt” and so any approach to their mapping has broader applicability. Of note is that the size of the brick kilns would be close to the pixel size of popular environmental remote sensing systems such as those carried on Landsat satellites, making them hard to observe, but they are considerably larger than the pixel size of fine spatial resolution satellite sensors and hence visually readily identifiable in such imagery, see Figure 1. Moreover, with resources such as Google Earth, free fine spatial resolution satellite sensor data are readily accessible. The following subsections summarise the data and methods used to explore the potential for contemporary methods and data sets to be applied to map brick kilns and contribute to the achievement of UN SDG 8.7.



(a)



(b)

Figure 1. Examples of brick kilns in the test site. (a) Oval and (b) circular. Note in (b) what appears to be an abandoned kiln.

2.1. Study Area and Imagery

Attention focused on an approximately 120 km² region in Rajasthan, India, see Figure 2. This region was selected as it is known to have a higher than average density of brick kilns [12] and so should furnish sufficient cases to both train and evaluate the classification methods. The machine learning methods were applied to the remotely sensed imagery displayed in Google Earth. The imagery, in RGB format, were downloaded for the analyses. The classifications were trained using imagery in the period 2003–2016 for the small region highlighted in the black box in Figure 2. The quality of the brick kiln classifications produced was evaluated using data from the broader study area, highlighted by the white box in Figure 2. For the accuracy assessment, attention focused on imagery of this region acquired in 2017 by the WorldView-2 system with a spatial resolution of approximately 0.5 m available in Google Earth. The region used for accuracy assessment contained 178 brick kilns, all identified by visual classification.




Figure 2. Imagery of the test site, to the east of Jaipur which is visible at the left-hand side of the Figure. The study focused on the region contained within the white box with the area inside the black box used only for the purposes of training the classifications. The image backdrop was acquired in January 2018.

2.2. Analyses

The brick kilns were mapped from the imagery contained in Google Earth using contemporary classifiers. Given that the brick kilns are distinct objects of characteristic appearance, the conventional classifiers commonly used in remote sensing research that are based mainly on spectral information were not used. Instead, attention focused on contemporary machine learning methods that have the potential to identify objects in images using especially, shape information. The basic nature of the methods is that they learn to identify the object of interest after presentation of a set of training examples. Key features of brick kilns to be learned relate to their size and shape as well as the presence of a central chimney that often casts a long shadow, see Figure 1, and need to be distinguished from other objects that may show similarity (e.g., road roundabouts) including abandoned kilns, see Figure 1b.

To meet the aims of the study, attention focused on three sets of analyses. First, the imagery were initially classified using the Faster R-CNN. Second, a reclassification of the outputs obtained from the Faster R-CNN was undertaken with a CNN. Third, the accuracy of the brick kiln classifications

obtained from the two previous analyses was assessed using human visual interpretations obtained by crowdsourcing as ground reference data.

2.2.1. Classification with Faster R-CNN

The Faster R-CNN was used to detect the brick kilns in the imagery. The Faster-RCNN is a two-stage object detector, including the Region Proposal Network (RPN) module and the Fast R-CNN module. RPN is used to generate candidate object regions, and the Fast R-CNN module is used to classify these candidate object regions and refine their locations. Both modules share the same convolutional features, and the VGG16 net, which contains 13 convolutional layers, was used in this study [28].

Training data comprised brick kilns identified within a small region, highlighted by the black box in Figure 2, which was studied in detail. These training data were extracted from Google Earth, using imagery of the study area acquired within the period 2003–2016. All images were divided into 800 × 800-pixel sub-images. Those sub-images that contained a brick kiln, identified visually, were used to form the training data set for the Faster R-CNN. Since the machine learning methods used typically require large training sets, replicates were produced by rotating the selected images by 90°, 180° and 270° to further increase the number of training samples available. In total, 572 training images were acquired, and they provided 1084 representations of brick kilns to inform the learning process. Examples of the training images used are shown in Figure 3.



Figure 3. Examples of training samples used in the Faster-RCNN analysis. The yellow rectangles identify kilns in the images. Image data shown from Google Earth provided by Digital Globe.

The Faster R-CNN was trained end-to-end by the Stochastic Gradient Descent (SGD) approach, using the training imagery along with annotated bounding boxes around each brick kiln. During the training of the classifier, a proposed region that has an Intersection-over-Union (IoU) overlap higher than 0.7 with any ground reference box of brick kilns was considered as a sample of a brick kiln, while that which had an IoU ratio lower than 0.3 was considered as a sample of background (i.e., a member of the non-kiln class). With the trained Faster R-CNN, brick kilns in the whole study area can then be detected. To overcome the graphics processing unit memory bottleneck, the detection over the whole study area was performed using smaller sliding windows in the large image. The original image was scanned with the window at the size of 800 × 800 pixels, and an overlap region, based on typical kiln size, of 200 × 200 pixels was used. The non-maximum suppression method was finally applied to eliminate the repeated detection in the overlap regions. In total, the training time was approximately 22 h.

The focus in this work is on the areas identified as potentially containing a brick kiln by the Faster R-CNN with a probability >0.5, although other probability thresholds were explored. The end product from the application of the Faster R-CNN to the imagery was an output in which the potential location of proposals, sites that could potentially contain a brick kiln, was highlighted.

2.2.2. Re-Classification with a CNN

The map produced from the Faster R-CNN analysis could be enhanced in a variety of ways. One is to refine the predictions with a further classification. Thus, a second analysis was undertaken based on the application of an additional classification achieved here by applying a CNN to the outputs that had been generated from the Faster R-CNN analysis.

The GoogLeNet was used as the CNN classifier [35]. The input to the CNN classifier was an image with the size of 221×221 pixels, and the output was the class that this image belonged to. Here, each input image was classified as either brick kiln or not. The same 572 images used to train the Faster R-CNN model were used to fine tune the pre-trained GoogLeNet-based CNN classifier. Training images were prepared for each class. For the class of kiln, the 1084 tagged representations of kilns were used. A 221×221 -pixel image in which the brick kiln is located in its centre was extracted for each tagged kiln, and all 1084 extracted images constitute the training samples for the class of kiln. The 572 small training images were again fed into the Faster R-CNN model, and all proposal regions for kilns with a probability >0.5 were extracted. Among those proposal regions, there are many background regions that were wrongly classified as brick kilns. These incorrectly classified regions, often called negative samples, were included as training samples of the non-kiln class. Moreover, some background regions that do not include kilns were randomly selected from these training images and also used as the training samples of the non-kiln class. Examples of training samples for the kiln and non-kiln class are shown in Figure 4. In total, the CNN training time was approximately 6 h.



Figure 4. Examples of training samples used in the convolution neural networks (CNN) analysis. The **top** row of images shows examples with brick kilns while the **bottom** row shows non-kiln samples. Image data shown from Google Earth provided by Digital Globe.

Once trained, the CNN classification was applied to refine the class labelling for the sites identified as potentially containing kilns by the Faster-RCNN model. For each proposal region detected as a brick kiln by the Faster R-CNN, a 221×221 -pixel image in which the centre of the proposal region is located in its centre was extracted. This image is further re-classified by the CNN classification. Images are either classified as kiln or non-kiln, and this process acts to reduce over-estimation of the kilns by the initial classification analysis by the Faster R-CNN.

2.2.3. Accuracy Assessment

The quality of the classifications obtained from the Faster R-CNN and that refined by the application of the CNN were assessed. This analysis was based on visual interpretation of the proposal regions highlighted in the outputs of the two machine learning methods. The image interpreters were those who had their labelling of brick kilns validated via a crowdsourcing analysis and hence were regarded as accurate labellers [12]. Similar to [12], each proposal region was labelled as kiln or non-kiln by human interpretation.

The quality of each of the brick kiln mappings obtained was estimated using standard measures of accuracy. A key focus was on the producer's accuracy or sensitivity of brick kiln mapping (the conditional probability of a proposal being labelled as a brick kiln and actually representing an area containing a brick kiln) and the errors of commission (cases that are non-kiln but mapped as kiln) and omission (cases of a brick kiln that were not mapped as such). Critical to the broader project in which this work is based is that errors of omission are regarded as being particularly severe; a map

that over-predicts brick kilns but rarely if at all omits kilns is of greater value to anti-slavery groups working on the ground.

3. Results and Discussion

The output of the Faster R-CNN is, essentially, an image with a set of proposals for brick kilns highlighted by rectangular bounding boxes, see Figure 5. Each proposal is predicted to contain a brick kiln with an estimated probability greater than a predefined threshold. For the test site used to evaluate classifier performance, see Figure 2, a series of output thresholds were explored. In trial analyses, when a high probability such as 0.8 was used, proposals were highly accurate but it was evident that many kilns were being omitted. With the 0.5 probability threshold, 366 brick kilns were predicted for the region used in evaluating classification outputs. This was a substantial over-estimate but, critically, was associated with zero omission error, see Table 1. With all of the 178 kilns located in the area used to evaluate classification accuracy included in the map output, the producer's accuracy for kiln mapping was 100%. Although the test site is small, this result does suggest a potential for machine learning as a tool for mapping brick kilns over the entire Brick Belt. The one drawback in the results was the large error of commission, with 188 cases that were not actually brick kilns classified as brick kilns. However, these latter errors could easily be reduced or even removed by adding a second classifier. For example, all 366 proposals could be presented to human interpreters and labelled. Although calling on human input, this demand is highly focussed, with the human interpreters looking at only a very small proportion of the area.

Table 1. Confusion matrix for the output of the analyses based on the Faster R-CNN. Columns show the labelling in the reference data while rows the labelling in the classifier output.

	Kiln	Non-Kiln	Σ
Kiln	178	188	366
Non-Kiln	0	0	0
Σ	178	188	366

The outputs of the Faster R-CNN analysis could also be fed into a second automated classifier. This was explored using a CNN as the second stage classifier. It was evident that the addition of a second automated classification to the original Faster R-CNN output was very effective at removing most of the commission errors that had been committed by the Faster R-CNN. Specifically, of the commission errors arising from the Faster R-CNN classification, all but nine were removed and labelled as non-kiln, see Table 2. A price paid for this, however, was an increase of omission errors to nine. The producer's accuracy for brick kilns was, therefore, 94.94%, but the classification had both low omission and commission errors. Again, visual classification could be used to further enhance the results if required. Critically, however, the accuracy of the classification arising from the CNN applied to the output of the Faster R-CNN analysis is such that users on the ground would waste little time at mislabelled sites.

Table 2. Confusion matrix for the output of the analyses based on the CNN applied to the predictions from the Faster R-CNN. Layout as in Table 1.

	Kiln	Non-Kiln	Σ
Kiln	169	9	178
Non-Kiln	9	179	188
Σ	178	188	366

In this study, some tolerance to error in the automated classification is provided, notably by the ability to usefully combine automated classification with highly targeted visual classification. A key feature in the analysis is the omission errors, sites of potential slavery activity that are missed and

hence would not attract the attention of anti-slavery work on the ground. Of particular note is that none of the brick kilns were omitted in the output of the Faster R-CNN. The large commission error, however, could lead to wasteful and inefficient anti-slavery activity on the ground. The use of a second classification, whether by manual interpretation or automated, could greatly reduce the commission errors while causing only small omission errors.




Figure 5. Extract of the output from the Faster R-CNN. Yellow bounding boxes show proposals of kilns. Within the extract shown all brick kilns present were identified with one commission error. Image data shown from Google Earth provided by Digital Globe.

Critically, the results indicate that contemporary machine learning methods may be used to accurately identify sites known to be strongly associated with slavery from remote sensing images and used to help achieve UN SDGs. The next stage of our work is to refine the approaches further but also to develop a wall-to-wall map of the entire Brick Belt in a quick and efficient way to aid anti-slavery actions on the ground. Potential avenues of research focus on refinements of the network structure used in the Faster R-CNN and the additional CNN classifier, as many novel network structures have been proposed recently. Other CNN-based object detection frameworks, such as the single shot multibox detector (SSD) [36] and you only look once (YOLO) [27], could also be used for comparison.

The parameterisation of these models should also be further explored for optimal classification. In addition, expansion of the training samples, perhaps including brick kilns with different shape and appearance, together with the use of more background regions that have a similarity with brick kilns could help enhance classification accuracy and generalizability. A clear potential for accurate mapping of brick kilns over large regions exists. Moreover, with the aid of the open Landsat data archive, we aim to explore the spatio-temporal pattern in brick kiln distribution to aid understanding of the brick manufacturing industry in the region. It must, however, be stressed that these approaches focus on a proxy variable for slavery by identifying kilns and not slaves directly, and hence additional work is also required to ensure that interventions on the ground recognise that some kilns mapped may use no slave labour at all. Finally, it should be stressed that the technologies used also have a real potential to be used in the detection of other activities that are known to use slavery (forced labour), for example, mining and illegal logging [37]. Further, the technologies also provide a means to study the associated environmental improvements that should arise from addressing slavery, enabling the freedom dividend to be characterised and quantified as well as contributing to other UN SDGs (for example UN SDGs 12 and 15).

4. Conclusions

Information on brick kilns in the Brick Belt is required to support anti-slavery activities. Given the large geographical extent of the Brick Belt, remote sensing has considerable potential as a source of data to inform efforts to detect and map brick kilns. Previous work has shown that visual interpretations obtained via crowdsourcing allow brick kilns to be accurately mapped from fine spatial resolution satellite sensor images such as those available in Google Earth. The ability to maintain a motivated crowd for accurate mapping of large areas is, however, a challenge and the time taken to produce the maps is a potential limitation. Here, the potential to use recent advances in key technologies for mapping brick kilns was explored.

Approaches for brick kiln detection that are fully or largely automated were shown to be able to produce accurate maps of brick kilns. Specifically, a fully automated approach using the Faster R-CNN yielded a map with a producer's accuracy for kilns of 100%. The high commission error associated with this classification could be a limitation, but this could be reduced by the application of a second classifier. Given that outputs of the initial classification allow activity to be focussed on a very small proportion of the region of study, visual classification and crowdsourcing should be feasible and able to manually adjust the outputs of the original classification. Alternatively, for a fully automated mapping approach, a further digital classifier may be applied. Here, it was shown that the application of a CNN classifier to the original outputs of the Faster R-CNN yielded a highly accurate map of brick kilns, with an overall accuracy of 95.08% and a producer's accuracy for brick kilns of 94.94% that had both low omission and commission errors. These results suggest considerable potential for mapping large areas quickly and accurately, which should help support anti-slavery activity on the ground and thereby aid the achievement of the relevant UN Sustainable Development Goals.

Author Contributions: D.S.B., G.M.F., F.L. and X.L. conceived and designed the research; D.S.B. and J.W. led ground data acquisition; F.L. and X.L. led machine learning analyses; G.M.F. led accuracy assessments. All authors were involved in interpretation. G.M.F. wrote the paper with inputs from all co-authors.

Funding: This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/R512849/1], awarded to G.M.F.; the British Academy Visiting Fellowships Programme, under the UK Government's Rutherford Fund, which supported X.L.'s work at the University of Nottingham with G.M.F. and D.S.B.; and the Rights Lab of the University of Nottingham.

Acknowledgments: The help of Manoj Arora (PEC University of Technology) and Malay Kumar (Prayatn Sanstha) as well as the provision of Google Earth is gratefully acknowledged. This publication uses data generated via the Zooniverse.org platform, development of which is funded by generous support, including a Global Impact Award from Google, and by a grant from the Alfred P. Sloan Foundation. We are grateful to the three referees for their constructive comments on the paper.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. GSI. *The Global Slavery Index 2016*, 3rd ed.; Walk Free Foundation: Nedlands, Australia, 2016.
2. United Nations. Sustainable Development Goals. Department of Economic and Social Affairs. Available online: <https://sustainabledevelopment.un.org/sdgs> (accessed on 7 November 2016).
3. United Nations. Sustainable Development Goal 8. Sustainable Development Knowledge Platform. Department of Economic and Social Affairs. 2016. Available online: <https://sustainabledevelopment.un.org/sdg8> (accessed on 7 November 2016).
4. Jackson, B.; Bales, K.; Owen, S.; Wardlaw, J.; Boyd, D.S. Analysing slavery through satellite technology: How remote sensing could revolutionise data collection to help end modern slavery. *J. Mod. Slavery* **2018**, *4*, 169–199. Available online: <http://slavefreetoday.org/journal/articlesissues/v4i2SEfullpub.pdf#page=180> (accessed on 28 January 2019).
5. Lechner, A.M.; McIntyre, N.; Raymond, C.M.; Witt, K.; Scott, M.; Rifkin, W. Challenges of integrated modelling in mining regions to address social, environmental and economic impacts. *Environ. Model. Softw.* **2017**, *268–281*. [CrossRef]
6. Bales, K. *Disposable People: New Slavery in the Global Economy*; University of California Press: Berkeley, CA, USA, 2012.
7. Kara, S. *Bonded Labor: Tackling the System of Slavery in South Asia*; Columbia University Press: New York, NY, USA, 2014.
8. Khan, A.; Qureshi, A.A. *Bonded Labour in Pakistan*; Oxford University Press: Oxford, UK, 2016.
9. ILO. *Unfree Labour in Pakistan: Work, Debt and Bondage in Brick Kilns*; International Labour Office: Geneva, Switzerland, 2005.
10. Save the Children. *The Small Hands of Slavery*; Save the Children Fund: London, UK, 2007.
11. Khan, A. *Over 250,000 Children Work in Brick Kilns*; The Express Tribune: Karachi, Pakistan, 3 October 2010; Available online: <https://tribune.com.pk/story/57855/over-250000-children-work-in-brick-kilns/> (accessed on 8 August 2017).
12. Boyd, D.S.; Jackson, B.; Wardlaw, J.; Foody, G.M.; Marsh, S.; Bales, K. Slavery from space: Demonstrating the role for satellite remote sensing to inform evidence-based action related to UN SDG number 8. *ISPRS J. Photogramm. Remote Sens.* **2018**, *142*, 380–388. [CrossRef]
13. Knoth, C.; Slimani, S.; Appel, M.; Pebesma, E. Combining automatic and manual image analysis in a web-mapping application for collaborative conflict damage assessment. *Appl. Geogr.* **2018**, *97*, 25–34. [CrossRef]
14. Gorelick, N.; Hancher, M.; Dixon, M.; Ilyushchenko, S.; Thau, D.; Moore, R. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sens. Environ.* **2017**, *202*, 18–27. [CrossRef]
15. A Fonte, C.C.; Bastin, L.; See, L.; Foody, G.; Lupia, F. Usability of VGI for validation of land cover maps. *Int. J. Geogr. Inf. Sci.* **2015**, *29*, 1269–1291. [CrossRef]
16. Foody, G.M. Citizen science in support of remote sensing research. In Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS), Milan, Italy, 26–31 July 2015; IEEE: Piscataway, NJ, USA; pp. 5387–5390.
17. Fritz, S.; See, L.; Perger, C.; McCallum, I.; Schill, C.; Schepaschenko, D.; Duerauer, M.; Karner, M.; Dresel, C.; Laso-Bayas, J.C.; et al. A global dataset of crowdsourced land cover and land use reference data. *Sci. Data* **2017**, *4*, 170075. [CrossRef] [PubMed]
18. See, L.; Mooney, P.; Foody, G.; Bastin, L.; Comber, A.; Estima, J.; Fritz, S.; Kerle, N.; Jiang, B.; Laakso, M.; et al. Crowdsourcing, citizen science or volunteered geographic information? The current state of crowdsourced geographic information. *ISPRS Int. J. Geo-Inf.* **2016**, *5*, 55. [CrossRef]
19. Goodchild, M.F. Citizens as sensors: The world of volunteered geography. *GeoJournal* **2007**, *69*, 211–221. [CrossRef]
20. Simpson, R.; Page, K.R.; De Roure, D. Zooniverse: Observing the world’s largest citizen science platform. In Proceedings of the 23rd International Conference on World Wide Web, Seoul, Korea, 7–11 April 2014; pp. 1049–1054.

21. Olteanu-Raimond, A.M.; Hart, G.; Foody, G.M.; Touya, G.; Kellenberger, T.; Demetriou, D. The scale of VGI in map production: A perspective on European National Mapping Agencies. *Trans. GIS* **2017**, *21*, 74–90. [[CrossRef](#)]
22. Peddle, D.R.; Foody, G.M.; Zhang, A.; Franklin, S.E.; LeDrew, E.F. Multi-source image classification II: An empirical comparison of evidential reasoning and neural network approaches. *Can. J. Remote Sens.* **1994**, *20*, 396–407. [[CrossRef](#)]
23. Mather, P.; Tso, B. *Classification Methods for Remotely Sensed Data*; CRC Press: Boca Raton, FL, USA, 2016.
24. Mountrakis, G.; Im, J.; Ogole, C. Support vector machines in remote sensing: A review. *ISPRS J. Photogramm. Remote Sens.* **2011**, *66*, 247–259. [[CrossRef](#)]
25. Ma, L.; Li, M.; Ma, X.; Cheng, L.; Du, P.; Liu, Y. A review of supervised object-based land-cover image classification. *ISPRS J. Photogramm. Remote Sens.* **2017**, *130*, 277–293. [[CrossRef](#)]
26. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*; The MIT Press: Cambridge, MA, USA, 2012; pp. 1097–1105.
27. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You only look once: Unified, real-time object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 779–788.
28. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *39*, 1137–1149. [[CrossRef](#)] [[PubMed](#)]
29. Yao, Y.; Jiang, Z.; Zhang, H.; Cai, B.; Meng, G.; Zuo, D. Chimney and condensing tower detection based on faster R-CNN in high resolution remote sensing images. In Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS), Fort Worth, TX, USA, 23–28 July 2017; IEEE: Piscataway, NJ, USA, 2017; pp. 3329–3332.
30. Deng, Z.; Sun, H.; Zhou, S.; Zhao, J.; Zou, H. Toward fast and accurate vehicle detection in aerial images using coupled region-based convolutional neural networks. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2017**, *10*, 3652–3664. [[CrossRef](#)]
31. Zhang, X.; Wang, Q.; Chen, G.; Dai, F.; Zhu, K.; Gong, Y.; Xie, Y. An object-based supervised classification framework for very-high-resolution remote sensing images using convolutional neural networks. *Remote Sens. Lett.* **2018**, *9*, 373–382. [[CrossRef](#)]
32. Cui, W.; Zheng, Z.; Zhou, Q.; Huang, J.; Yuan, Y. Application of a parallel spectral–spatial convolution neural network in object-oriented remote sensing land use classification. *Remote Sens. Lett.* **2018**, *9*, 334–342. [[CrossRef](#)]
33. Chen, F.; Ren, R.; Van de Voorde, T.; Xu, W.; Zhou, G.; Zhou, Y. Fast automatic airport detection in remote sensing images using convolutional neural networks. *Remote Sens.* **2018**, *10*, 443. [[CrossRef](#)]
34. Ren, Y.; Zhu, C.; Xiao, S. Deformable Faster R-CNN with aggregating multi-layer features for partially occluded object detection in optical remote sensing images. *Remote Sens.* **2018**, *10*, 1470. [[CrossRef](#)]
35. Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; Rabinovich, A. Going deeper with convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Boston, MA, USA, 7–12 June 2015; IEEE: Piscataway, NJ, USA, 2015; pp. 1–9.
36. Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.Y.; Berg, A.C. SSD: Single Shot MultiBox Detector. In *European Conference on Computer Vision*; Springer International Publishing: Cham, Switzerland, 2016; pp. 21–37.
37. Bales, K. *Blood and Earth: Modern Slavery, Ecocide, and the Secret to Saving the World*; Spiegel & Grau: New York, NY, USA, 2016; 304p.

