

Article

Human Lights

Ilari Määttä ^{1,*}  and Christian Lessmann ^{1,2}¹ Technische Universität Braunschweig, 38106 Braunschweig, Germany; c.lessmann@tu-braunschweig.de² Ifo Institute and CESifo, 81679 Munich, Germany

* Correspondence: i.maaettae@tu-braunschweig.de

Received: 25 July 2019; Accepted: 17 September 2019; Published: 20 September 2019



Abstract: Satellite nighttime light data open new opportunities for economic research. The data are objective and suitable for the study of regions at various territorial levels. Given the lack of reliable official data, nightlights are often a proxy for economic activity, particularly in developing countries. However, the commonly used product, Stable Lights, has difficulty separating background noise from economic activity at lower levels of light intensity. The problem is rooted in the aim of separating transient light from stable lights, even though light from economic activity can also be transient. We propose an alternative filtering process that aims to identify lights emitted by human beings. We train a machine learning algorithm to learn light patterns in and outside built-up areas using Global Human Settlements Layer (GHSL) data. Based on predicted probabilities, we include lights in those places with a high likelihood of being man-made. We show that using regional light characteristics in the process increases the accuracy of predictions at the cost of introducing a mechanical spatial correlation. We create two alternative products as proxies of economic activity. Global Human Lights minimizes the bias from using regional information, while Local Human Lights maximizes accuracy. The latter shows that we can improve the detection of human-generated light, especially in Africa.

Keywords: nighttime lights; human light sources; economic activity; human settlements; machine learning; random forest; DMSP-OLS; GHSL

1. Introduction

Nighttime lights detected in satellite images of the Earth have become a widely used measure for approximating national and regional economic activity. Early studies that established nighttime lights as a proxy of economic activity include References [1–5], the fundamental research of which opened new opportunities for economic research at the regional level on topics such as geography and trade [6,7], institutions [8], inequality [9,10], regional favoritism [11], natural resources and conflict [12], and the impact of natural disasters [13], among others. The broad use of nighttime light data demonstrates how valuable it is in current economic research.

Light data hold two major advantages when compared to more established measures of development, such as gross domestic product (GDP). First, it is independent of national or regional borders, which gives researchers the freedom to define the units of interest. Second, the data are collected in a uniform fashion for the whole world, which solves issues of data quality across countries. This is particularly important for developing countries, where economic statistics are less reliable compared to the developed world [1]. These advantages contribute to the recent popularity of light data in the field of economics in general and development economics and regional science in particular.

The National Geophysical Data Center (NGDC) at the National Oceanic and Atmospheric Administrations (NOAA) provides nighttime light data. Researchers process the satellite images collected by the United States (US) Air Force Weather Agency (AFWA) Defense Meteorological Satellite Program (DMSP) into easy-to-use and freely available products. Stable Lights is the most widely used

product in economics [14]. It removes ephemeral lights (e.g., forest fires) and background noise from global composites of annual average light. The resulting product identifies areas of persistent lighting, which is a reasonable proxy for economic activity.

The wide use of Stable Lights in economics has put the data-generating process under increased scrutiny. Even though more recent and precise data are already available (e.g., the Visible Infrared Imaging Radiometer Suite, VIIRS), perfecting nighttime light data is still important for economic research. Economic development takes time, and therefore the long historical data provided by the Defense Meteorological Satellite Program's Operational Linescan System (DMSP-OLS) sensors is invaluable for the field. The interest in improving the Stable Light product has generated solutions to problems such as "top-coding" [15] and blurring [16]. However, another important issue is still understudied. Bruederle and Hodler [17] analyzed the relationship between Stable Lights and measures of human development based on the Demographic Health Surveys (DHS). They show that nighttime lights are positively associated with location-specific indicators of human development. However, they also point out a problem at the lower tail of the light distribution: half of the survey locations marked in 29 African countries under study do not emit any light. Nevertheless, the authors assume that these inhabited locations do emit light, but the emissions might not be strong enough to be recorded by the satellites, or they are filtered out as background noise in the Stable Lights process. This issue, which is particularly important for development economics research, is the starting point for our study. We investigate whether the nightlights of inhabited locations at the lower tail of the light distribution are indistinguishable from background noise, or if the Stable Lights process is too strict in filtering them out.

To the best of our knowledge, the problem of Stable Lights filtering out low-coded but inhabited pixels has not been addressed so far. Therefore, we contribute to the literature by addressing this "bottom-coding" problem. We find that the Stable Lights process is too strict and filters out light from human activity, which is not just background noise. Possible explanations may be, for example, that electricity might not be available every day or that economic activity is seasonal. Such light emissions may be classified as transient in the Stable Lights process and left out from the final product. As a result, lights from cities as large as 100,000 inhabitants are removed in the Stable Light filtering process. This problem emerges particularly in developing countries, where the nighttime light data has the most potential due to lacking quality of other data sources. We argue that even though the Stable Lights process is efficient in separating transient light from stable light, it is suboptimal in separating between human-generated light and light that occurs without human action. This causes a problem in economic research, where the Stable Lights product is used as a proxy for economic activity. To properly measure economic activity, it is necessary to keep the light also from transient light sources, if it is generated by economic activity and distinguishable from background noise. Hence, our goal differs from Stable Lights by aiming to separate background noise from human-generated light.

Pixel classification into human-generated light and background noise is difficult, because the level of the background noise differs across regions. If we relax the filtering process and add more light, we also detect more human-generated light. However, we make correct detections with diminishing returns because we misclassify an increasing number of background noise pixels once we approach the lower tail of the light distribution. In other words, it is easier to separate background noise from big city lights than villages with dim light emissions. The optimal amount of light in a final product depends on the particular research application at hand. However, the order in which we should classify the pixels as lit can be determined with a classification rule. Machine learning methods perform very well on classification tasks in terms of out-of-sample prediction accuracy. We use the random forest algorithm to learn patterns in the light data that are associated with human-generated light. The light data include the average visible band and frequency of light detection images from NOAA and regional aggregates derived from them. Human-generated light is proxied by built-up content layer from the Global Human Settlements project (GHSL). It labels the areas where the light patterns are potentially of human origin. The GHSL built-up grid searches daylight Landsat satellite images

for man-made objects and records the proportion of the building footprint area within a pixel [18]. The idea of the process is that we can learn what kind of light patterns are emitted in built-up areas and then predict the probability of light being from a built-up area by only using the light characteristics (not the built-up data itself). We use these predicted probabilities to add the light in the order that most likely chooses human-generated light over background noise.

Moreover, we find three additional problems in the Stable Lights data that should be recognized in economic research: (1) the background light noise filtering process introduces measurement error that is correlated with regional light characteristics. This is a concern for studies that rely on the assumption of uncorrelated measurement errors in the light data and a variable of interest such as GDP [5]; (2) the filtering process adds spatial correlation in the light data. Our concern is that studies on the diffusion of economic activity may find an effect that partially emerges in the filtering process instead of the true spatial correlation in economic activity [17]; (3) the reproduction of the Stable Lights process is impossible with the data that is freely available. Any process to remove the background noise in nightlight data requires some assumptions of characteristics that separate background noise from the relevant signal. In the Stable Lights process, these assumptions are made by basing the level of local background noise on hand-drawn light-free areas. The underlying assumptions of the filtering process have a direct impact on results in economic research, which is why it is crucial to be able to evaluate their impact.

The different issues mentioned above cause trade-offs depending on what problems should be solved most effectively. Therefore, it is not possible to maximize the accuracy in filtering out the background noise without making compromises in the three above-mentioned problems. We provide two different solutions which future economic research might consider. The Global Human Lights product addresses the concerns of correlated measurement error and spatial correlation. It runs the same process for the whole world and uses only the most necessary local information to minimize regional differences in the filtering rule. However, the key to improved precision in separating human-generated light from background noise relies on that local information. Therefore, the Local Human Lights product takes full advantage of the regional light characteristics to maximize the efficiency in separating human-generated light from the background noise. This leads to a significant increase in the accuracy of detection of human-generated light but comes at the cost of introducing regional bias. Fortunately, appropriate econometric methods can address the bias, and we suggest the selection of particular settings of our method depending on the research application. Both products are intended for use in research applications that use nighttime lights as a proxy for economic activity. Together with Stable Lights, the Human Lights products can be used to assess how an alternative filtering process impacts the results of the research application at hand. Global studies on economic activity might consider the Global Human Lights data, while regional country-case studies might take advantage of the characteristics of Local Human Lights.

Our results show that the Stable Lights product is too strict in filtering out the background noise, and therefore leaves out a considerable amount of light from economic activity at the lower tail of the light distribution. Furthermore, the missing light from economic activity is not equally distributed across regions. We use our Human Lights products to show that the detection of light from economic activity can be improved, especially in developing countries. We also demonstrate that it is important to be transparent and evaluate the background noise filtering process because it has a direct impact on the results of research applications. Our study considers satellite images taken by the satellite F15 in the year 2001. We chose the satellite/year randomly between the six options in the middle of the nighttime light data time frame (years 2000–2002, satellites F14 and F15). We limited the choice to these years in order to avoid historically low or high light levels at the start and end of the timeframe. The extreme cases would not have served as good examples of the Human Lights filtering process which we intend to extend to all satellites/years in the future. Our results are promising. With regard to the Local Human Lights product, we are able to improve the detection rates of settlement points in Africa recorded by the Global Rural–Urban Mapping Project (GRUMP) by about 15%.

Section 2 summarizes the relevant literature on the Stable Lights product history and its use in economics. Section 3 introduces our data and method in detail. Section 4 shows the results by comparing the Human Lights products to Stable Lights. Section 5 discusses the implications for economic research. Finally, Section 6 concludes the paper.

2. Nighttime Lights and Economic Activity

Since the 1970s, the U.S. Air Force has used DMSP-OLS to collect imagery of the Earth [19]. The images taken at night are designed to detect clouds by using a photomultiplier tube to intensify the visible and near-infrared signal (see Reference [16] for a detailed description of the imaging process). As an unintended consequence, the intensified light signal was also able to detect city lights, gas flares and forest fires, as first noted by Croft [20]. He discovered that lights of cities are clearly visible in the images, and he remarked upon their potential for socio-economic research.

Elvidge et al. [19] proposed the first method to separate transient light sources from stable light sources for mapping city lights. They carefully prepared the raw satellite data into a panel dataset and then used an algorithm that considered light intensity and frequency to isolate stable light from transient light. The final maps showed frequencies of light detection for each cell. By choosing a threshold, the authors could provide a binary map of stable lights. The Stable Lights process was further developed after the initial method. Baugh et al. [14] describe the current version of the Stable Lights process. We recap this process in detail to provide a well-informed comparison between the Stable Lights and our new Human Lights products.

Initially, the Stable Lights process re-projects the raw satellite images and removes those observations affected by brightness from the lunar cycle, summer months with late sunset, auroral activity and cloud cover. Four products result out of the data cleaning:

1. An average visible band image that shows average annual light values for all pixels;
2. A cloud-free coverage image that counts the total number of uncontaminated observations;
3. The frequency of light detection, which shows the percentage of days with the detection of light;
4. A histogram for each pixel, which shows the number of days for different light values.

Products 1–3 are all .tif format raster images which are available online for the years 1992–2013. The images are in Platte Carree projection and have a resolution of 30 arc seconds, meaning that the size of one pixel is roughly one square kilometer at the equator.

The goal of the Stable Lights process is to improve the average visible band image, because it contains background noise and ephemeral light sources such as forest fires [14]. Due to the background noise, the light intensity in a cell is never zero, but varies between values 2 and 63. The Stable Lights process identifies pixels that have outlier light values over time by using the histograms of the daily pixel light values. The intuition is that a pixel is located on a stable light source if the light intensity does not vary much during a year. If a pixel has only a few days of high light intensity, it is likely due to a transient light source such as forest fires or fishing boats. The authors define a rule-based algorithm to decide which observations of the daily pixel light values are outliers. After the outlier removal, the procedure exchanges the new pixel values into the average visible band image.

Baugh et al. [14] then proceeded to remove the background noise from the outlier-removed average visible band image. Their method starts with an analyst marking areas that appear “light-free” in the outlier-removed average visible band image by visual inspection. Then, they move a window of 400×400 pixels (a tile) 256 times around a smaller window of 25×25 pixels (a kernel) in 25-pixel steps horizontally and vertically for every kernel in the image. At every step, they record whether the light values of the pixels in the current kernel are higher than the maximum “light-free” values in the current tile. As a result, every pixel has 256 records indicating whether it is higher than the “light-free” areas around it. Finally, they create the Stable Lights image by leaving light values from the average visible band (not the outlier removed) for every pixel that is higher than “light-free” in at least in 40% of the 256 records; the values of rest of the pixels are set to zero.

The resulting annual Stable Lights composite images are widely used in economics. Henderson et al. [5] demonstrate that lights are a suitable proxy for economic activity at the national level. They start by noting that GDP is the most important indicator of economic development. However, they add that it suffers from conceptual problems of definition and measurement problems, especially in developing countries, and inconsistent availability at the sub-national level. The nightlight data remedy these problems by offering a consistent definition and measurement across all countries and the freedom to choose the area of interest irrespective of administrative borders. These properties opened opportunities for a wide array of research applications in economics, and the Stable Lights product was rapidly accepted as a valid proxy for economic activity.

The first advantage—a universal definition and measurement—was utilized by Ghosh et al. [21], who showed that the size of Mexico’s economy may be 150 percent larger than official Gross National Income (GNI) estimates due to the informal economy and remittances. Martinez [22] shows that authoritarian regimes inflate their yearly GDP growth rates disproportionately compared with democratic regimes. An example of the second advantage—being independent of administrative borders—facilitated a study by Alesina et al. [9] who constructed measures of inequality and development based on ethnic homelands. Data on economic activity for different ethnic groups within a country would be very difficult to gather on a large scale. However, nightlight data is easy to aggregate on any region that a researcher chooses. They found a strong negative association between ethnic income inequality and development.

The increased use of the Stable Lights product has attracted interest into its properties. Bluhm and Krause [15] addressed the problem known as “top-coding”, which means that the light intensity in a pixel is saturated at the value of 63. The top-coding problem complicates comparisons between cities and leads to an underestimation of income differences between rural and urban areas. Another issue is blurring, for which Abrahams et al. [16] provided a correction method. The problem is that light from strong sources blooms over neighboring pixels due to data collection and processing.

Our study addresses two further issues that are important for economic research: “bottom-coding” and background noise filtering. The need to address these problems is clear because they have a direct impact on the results of research applications. This is particularly true for rural, less developed areas, where nightlights have the largest potential, since other valid statistical data are usually unavailable or of insufficient quality.

3. Empirical Analysis

This section explains the data and methods we apply in our Human Lights products (Section 3.1). Since the amount of regional information used (Section 3.2) and the amount of light added (Section 3.3) are decisive in the process and influence the process design itself (Sections 3.4 and 3.5), we include a detailed discussion of these. We also include an example of the differences in the filtering process using Nigeria as a smaller region to demonstrate how results are affected by different choices (Section 3.6). Note that we always study the whole world, but we refer to different regions to explain our findings.

3.1. Method and Data

Making a distinction between light from human activity and light from other sources lies at the core of our paper. As exemplified by the Stable Lights product, one path would be to rely on analysis, visual inspection and experience to create decision rules. However, even the top experts cannot be certain if they have given the most relevant weights to different variables and their interactions. Fortunately, machine learning methods excel in doing just that. Supervised learning methods are capable of constructing a decision rule that optimally adjusts the weights of input variables. The algorithm finds the best model by trying out a large amount of candidate models, while judging their predictive performance out of the sample to avoid overfitting. We choose the random forest machine learning method because it is relatively easy to use and understand, even for those who have no experience

in machine learning methods. Furthermore, it performs well with nonlinear patterns in data and relatively well compared to more advanced machine learning methods [23].

The random forest algorithm requires training data in which correct outcomes are labeled. The training data allow the random forest to learn relevant patterns in input data that are associated with the outcomes. In the case of nightlights, we do not have direct data or a clear definition on which lights are from economic activity. Therefore, we rely on areas with human built-up content as a proxy for areas where lights are human-generated. Our major assumption is that human-generated light outside and inside built-up areas follows similar light patterns. If the patterns are similar, then we are able to detect human-generated light with matching patterns also in non-built areas. However, if there are differing human-generated light patterns in non-built areas, then those patterns will not be recognized as human-generated by the algorithm. Our assumption may not always hold for every human light source; e.g., for fishing that does not take place in a built-up area. Therefore, while our assumption might be violated in some cases, we believe that these cases are marginal and do not have a strong influence on the results. Despite this constraint, we think that built-up data comprise the best available choice to inform us about human-generated light characteristics on a global scale.

The outcome variable used in our model is the Global Human Settlement Layer (GHSL) built-up grid. It is based on Landsat satellite images and prepared following the GHSL methodology [18]. In short, it searches daylight images for man-made objects and records the proportion of the building footprint area within a pixel. The pixels in the resulting raster image take a value between 0% and 100%. We use the data at a 1 km² resolution per pixel for the years 1990, 2000 and 2014. We re-project and resample the data to align with the nightlight data and match the coordinate reference systems. We also interpolate the data between the years to match the period of the nightlight data. Finally, the built-up raster images are changed into binary format because our goal is to classify the light in pixels as human-generated or not. We code pixels containing built content as one and set other pixels to zero. It is not clear if we should expect human-generated light from pixels with less than 1% built-up content, for example; however, we do not want to miss any signal of potentially human-generated light by choosing an arbitrary threshold. Therefore, we include all pixels with built-up content.

We use several variables for light characteristics in each pixel and surrounding areas. The random forest model uses them to establish patterns in light that are associated with the built-up content outcome variable. The most important variables are the average visible band (*avg_light*) and frequency of light detection (*freq_light*) images from NOAA. They provide information on average light over one year and the percentage of days with light detections in a pixel.

In addition, we generate variables that describe light characteristics around a pixel. The regional input variables address concerns that the background noise in light data differs across regions. We use a *local_noise* variable to identify regions, where the background noise in *avg_light* is systematically higher. The variable counts the number of pixels below the value 6 in *avg_light* image in a square window of 499 × 499 pixels for all pixels. We pick the value 6 as a threshold because it is at the upper end of the background noise values. The variable informs the algorithm about the level of background noise in a region. For example, a pixel with a light value 7 is more likely to be of human origin in an area with low background noise than in an area where the noise is regularly over the value 6. Figure A1 in the appendix illustrates the *local_noise* variable. A quick glance shows that some disturbance from auroral activity remains in the *avg_light* image.

In a similar manner, we create a measure of light detection in the region around a pixel. It counts pixels with zero light detections in a 399 × 399 pixel region for all pixels. The *local_detections* variable helps to account for systematic regional differences in the *freq_light* variable. Figure A2 in the appendix shows the *local_detections* variable. We observe that the *freq_light* variable is influenced by unexplained regional differences. Especially in the area around Brazil, we observe light detections in almost every pixel, even over water. In contrast, northern parts have regions without any light detections. We do not know the exact reason for these differences; they may have been introduced in the earlier phase of the raw data cleaning process.

For the Local Human Lights product, we add more variables on local light characteristics to maximize the classification accuracy. We create these variables by taking averages of *avg_light* and *freq_light* on different area sizes around a pixel. For example, *lm_avg_199* calculates the average of *avg_light* in a 199×199 pixel area around a pixel. Similarly, we create the variables *lm_freq_5*, *lm_avg_25* and *lm_freq_99*. The reason for adding local averages on different area sizes is that we do not know the correct radius of spatial correlation in human-generated light. By adding multiple variables, we let the algorithm learn the best combination of the spatial information. An overview of the input variables used in each Human Lights product is provided in Table 1. We are only using input variables derived from the light products because we want the final products to remain independent from other data sources. This allows researchers to apply the data in combination with other data sources. This also distinguishes Human Lights from products that use any available variables to improve their estimates of economic activity.

Table 1. Human Lights input variables.

Input Variables	Global Human Lights	Local Human Lights
<i>avg_light</i>	X	X
<i>freq_light</i>	X	X
<i>local_noise</i>	X	X
<i>local_detections</i>	X	X
<i>lm_freq_5</i>		X
<i>lm_avg_25</i>		X
<i>lm_freq_99</i>		X
<i>lm_avg_199</i>		X

We evaluate the success of finding human-generated light by using built-up data and data from the Global Rural-Urban Mapping Project (GRUMP version 1, CIESIN 2011). An advantage of the built-up data is that it offers a measure of the classification success that considers all pixels. However, our metric of success is flawed if we evaluate the products based on how often light in (non-) built pixels is classified as (non-) human-generated. As mentioned above, human-generated light does not perfectly overlap with the built areas, which is why this evaluation approach is only a crude indicator of success. As an alternative metric of human-generated light detection, we use the GRUMP data on human settlements. These provide point data of human settlements and estimated populations. Unlike the GRUMP urban area extent dataset, the list of settlement points does not rely on nightlight data. We note that the points are not always exactly on top of the most densely populated area in a settlement. Therefore, we count a settlement as successfully detected if we find light in 3×3 pixel area around a settlement point.

3.2. Regional Information

The accuracy in classifying light in pixels as human-generated or background noise depends on the amount of regional information that we include in the filtering process. However, there is a tradeoff between the increased accuracy and the regional bias introduced by adding regional information. The bias results from an increasingly differential treatment of regions based on their light characteristics. For example, the hand-drawn light-free area thresholds in the Stable Lights process may be higher in highly lit areas, which makes it less likely for pixels in those areas to be classified as stable light. However, this is only speculation, because we cannot evaluate the magnitude or even the direction of the bias without knowledge of the location of light-free areas.

In economic research, the regional bias has two important implications: first, it generates a measurement error that is correlated with the regional light characteristics. For example, it is relatively easy to classify the pixels in Germany, where populated areas are more likely to emit a clear signal in the light data. In contrast, the classification is more difficult if a populated village and unpopulated forest in Nigeria have an indistinguishable light signal. Populated regions that have less overall

light emissions tend to have more pixels that are difficult to classify. If we vary the classification process between Nigeria and Germany, then the classification errors will depend on that variation as well. This leads to a problem in research applications that rely on the assumption of uncorrelated measurement errors in light data and target data (e.g., GDP).

Second, regional bias introduces spatial correlation. When we use regional information in the classification process, a pixel close to a highly lit area is more likely to be classified as human-generated than a pixel with the same *avg_light* and *freq_light* in a dark area. For the same reason as discussed above, we do not know how the spatial correlation bias affects the Stable Lights product. The spatial correlation of population and economic activity is not a new finding, but we emphasize that it is amplified in the nightlight data to a degree that depends on the amount of regional light information included in the classification process. Therefore, it has a direct impact on those research applications that study how economic shocks spread over space.

Of course, some regional information in the filtering process is necessary, because the background noise also differs across regions. If we do not consider these differences, then we would not be able to distinguish background noise from human-generated light at lower light levels. For example, a village with low light emissions may be detected only if the background noise in the region is low. Another village with similar light characteristics in a high background noise region might not be distinguishable. If we do not use regional information, then we would have to set a strict background noise filter that filters out both of the villages. In contrast, adding information to the regional background noise level allows us to detect the village in a low background noise region, thereby leading to better overall detection results. Therefore, we need to find a rule that gives us the optimal amount of regional information that minimizes the impact of both regional background noise and regional bias. Unfortunately, we cannot perfectly solve this task, since we would need to know the ground truth of human activities in every pixel and every year. Lacking such data, we have to assess the impact of different approaches instead. Such an assessment is only possible if the filtering process is transparent and reproducible. In the case of Stable Lights, we do not know the degree of regional information used in the process because the light-free areas are based on subjective choice.

We provide two benchmark Human Lights products which we use to assess the impact of regional information on research results. The Global Human Lights product uses only a minimal amount of regional information (*local_noise* and *local_detections*) and applies the same filtering rule for the whole world. The Local Human Lights product includes many additional regional variables and varies the filtering rule across regions. The degree of regional information in the Stable Lights filtering process is most likely somewhere between our products. Together, all three products can help researchers to assess the impact of varying the degree of regional information in the filtering process on research results at hand. Based on the case of Nigeria, we provide a concrete example for the impact of varying regional information in the filtering process (Section 3.6).

3.3. The Amount of Light

As with the amount of regional information, we do not know the exact correct number of pixels we should classify as human-generated light. Indeed, the optimal amount even depends on the specific research application. In some cases, researchers may prefer to find even the smallest signal of human-generated light, while in others it is important to minimize the number of misclassified pixels.

The number of pixels lit in Stable Lights is a good starting point. However, we do not know if the amount chosen in the Stable Light process is optimal. The threshold in Stable Lights, which determines the number of lit pixels, is that an outlier cleaned pixel must be above the local light-free values 40% of the time as the comparison window is moved around it. Because the light-free values differ across regions, the light amount is ultimately also set with a different threshold for every region. Unfortunately, there is no further explanation for the 40% threshold in Reference [14]. To construct the Human Lights products, we also need an arbitrary choice about the threshold, but we provide a description of how it is set and instructions on how to adjust it if required.

The random forest algorithm, which learns the human-generated light patterns from built-up pixels, provides us with classification probabilities. Those probabilities inform us about the right order to classify the pixels as lit. However, it does not give us information to select a probability threshold that determines the total amount of lit pixels. Usually, machine learning classification tasks rely on some measure of prediction success to determine the optimal threshold (e.g., the accuracy or F1 score). However, our outcome variable, the built-up data, is only a proxy for the true human-generated light. If the human-generated light would perfectly overlap with built areas, then we could simply use the built-up areas to filter out the background noise. However, humans generate light also in non-built areas (e.g., when fishing), and not all built areas emit detectable light (e.g., remote roads). In addition, if we choose the number of lit pixels based on prediction success, then we effectively tie the number of lit pixels to the number of built pixels. This would be problematic if researchers use the data to compare changes in light over time because we interpolate the built-up data. As a result, we would enforce an arbitrarily linear trend in light over time.

Our solution is to derive the amount of light from the average visible band image. In the *avg_light* image, pixels with a light value below 5 are most likely background noise. However, some light emissions may be of human origin. We decide how many of those low-light pixels we tolerate being classified as human-generated light. In other words, we keep lighting up pixels in the order of the highest probability of being human-generated, until pixels with the lowest average light values start lighting up. Adjusting the tolerance threshold indirectly sets the total amount of lit pixels in the Human Lights products. The advantage of our approach is that it can be applied over time and to satellites without having to make subjective choices about the light amount. The study regions will be equally affected by the changes in overall light intensity, as long as the tolerance threshold is kept constant. We will discuss the number of lit pixels further in the following subsections, which describe the specific steps to create the Global and Local Human Lights products.

3.4. Global Human Lights Process

Our previous discussion implies that the use of regional information in the process and the number of lit pixels have important effects on our results. Therefore, we provide two different products we can use together with the Stable Lights to assess how the filtering process settings affect the results. We suggest that researchers tailor the filtering process settings for the particular application, although re-doing the whole process for every study can be burdensome. We provide detailed instructions on how to reproduce the process and how to change the settings in supplementary materials. For those researchers who are content with our settings, we provide ready-to-use benchmark products.

We first introduce the Global Human Lights product. It minimizes the use of regional information by using the same classification model and threshold for the whole world. This approach mitigates the regional biases arising from correlated measurement errors and spatial correlation. However, we cannot perfectly eliminate the biases, because we have to add a bare minimum of regional information to account for regional anomalies (variables *local_noise* and *local_detections*). A global model without any regional information would result in classifying auroral activity and a relatively high amount of background noise as human-generated light. It would therefore bias the results more than the regional bias introduced in our chosen model. Figure 1 gives an overview of the Global Human Lights filtering process.

The process for the Global Human Lights product consists of four steps:

1. We prepare the input variables (*avg_light*, *freq_light*, *local_noise* and *local_detections*) and the outcome labels (*built_up_binary*) as explained in Section 3.1. The raster images are then stacked together and turned into a data table which the random forest algorithm can use. Observations in the data table correspond to pixels in the global raster images, which means that the table has more than 679 million rows.
2. We lack the processing power to run the random forest algorithm on such a huge dataset. However, we can take repeated random samples from the data and then average the predicted probabilities

of light in pixels being human-generated. We take random samples of 1% of the data and repeat the process 100 times. The predictions between samples do not differ significantly. Therefore, we are confident that our results are not driven by random variation in the repeated samples.

3. In Step 3, we set our tolerance of low average value pixels (below 5) classified as human-generated to 0.08 percent. We fix the tolerance level over time to avoid a subjective choice about the number of lit pixels for every year. We can then derive a probability threshold and mask pixels whose predicted probabilities jump over the threshold.
4. Finally, we use the light values from the average visible band image for the masked pixels and set other pixels to zero.

We provide descriptive statistics and discuss the use of the final Global Human Lights product in Section 4.

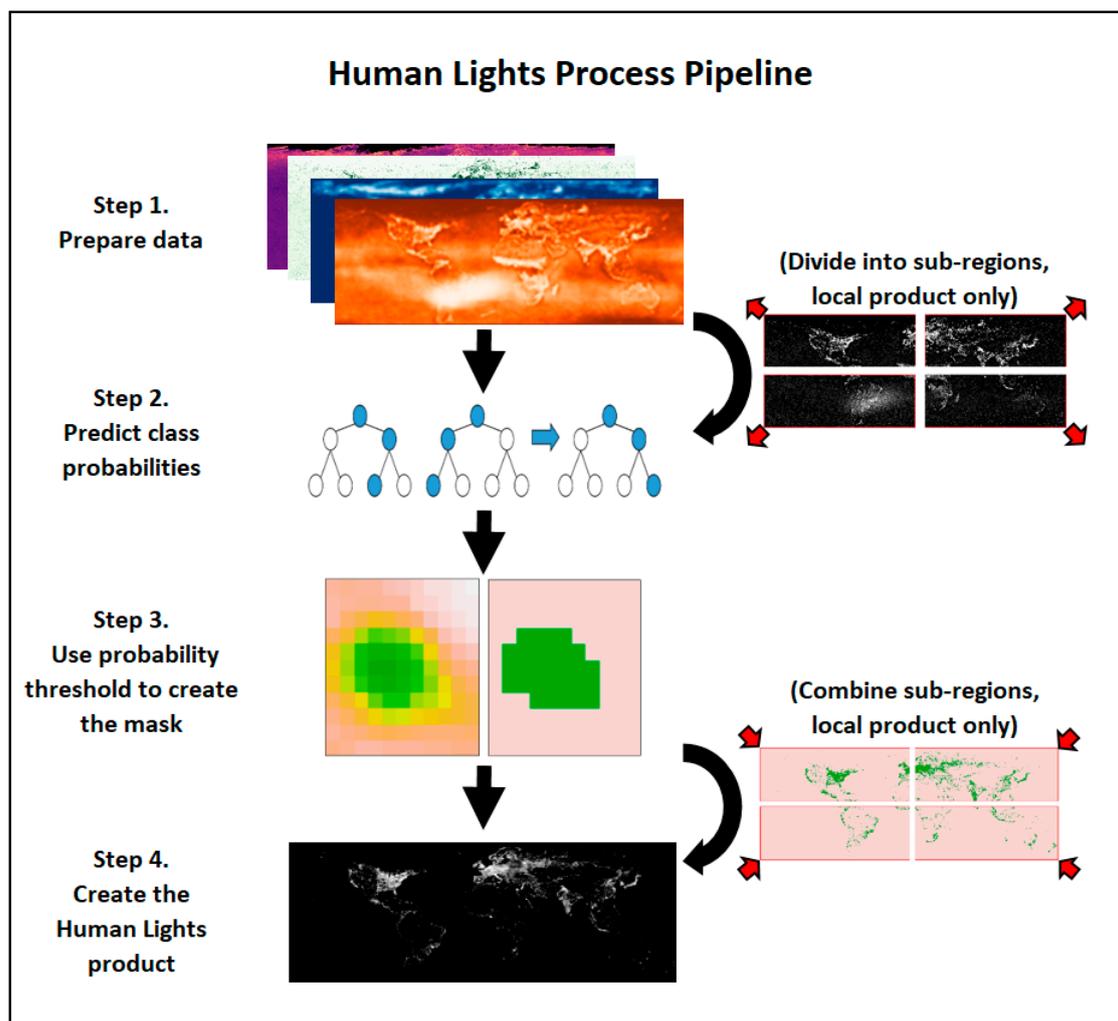


Figure 1. The Human Lights process.

3.5. Local Human Lights Process

Our second product shifts weight from minimizing regional bias to maximizing the accuracy in human-generated light detection. We utilize the regional light characteristics in the filtering process in two ways. First, we include more regional variables as inputs in the random forest algorithm. Second, we divide the world into sub-regions and run the algorithm separately for each region. This allows the algorithm to learn how the light characteristics relate to built-up areas in different sub-regions.

Since that relationship most likely differs across the world, we achieve a better prediction accuracy by allowing for variations in the classification rule.

In practice, we crop the variable images into 2000×2000 pixel sub-regions after the data preparation phase (Step 1). Accordingly, the size of a sub-region is approximately 3.4 million km² at the equator, which is close to the size of India. The number of rows in a sub-region is 4 million, and we take a subsample of 10% for the model training. By running different versions, we found this window size to have a good balance between improvements in prediction accuracy and being widely independent of the quality of the built-up data. However, the changes in the classification rule across the sub-regions causes a problem at their borders. The shifts in predicted probabilities, and consequently the number of lit pixels, are visibly clear.

The solution for smoothing the border effects is to move the sub-region window in smaller steps and then average the results. The step size should be as small as possible, but we rapidly run into computational constraints. Therefore, we decide to move the window in steps of 250 pixels so that we receive 64 values for each pixel. Unlike in the global approach, we now set the number of lit pixels within each window by using a 4% tolerance level. This means that we keep classifying light in pixels as human-generated in the order of highest probability until 4% of *avg_light* pixels with values below 5 are included. By setting the amount of light with this tolerance threshold, we avoid having to make subjective choices of light amounts across regions or base them on the built-up data. However, we have to include a backstop or we end up adding lit pixels also in areas where there is no human-generated light at all. Therefore, we additionally set all pixels with a predicted probability less than 20% of being human-generated to zero.

As described above, we classify light in each pixel 64 times by the overlapping sub-region windows as human-generated or not. After Step 3 in the filtering process, we mosaic all the windows and count the classifications. We then create the final product in Step 4 by taking the light value from the average visible band image if a pixel is classified as human-generated more than eight times out of 64. Otherwise, we set the light value in a pixel to zero. The requirement of eight human-generated light classifications makes a final global adjustment to the amount of light. We have chosen the light amount settings in a way that balances the improvement in detecting human-generated light while keeping misclassifications at a low level. However, the choices are arbitrary, and the optimal number of lit pixels depends on the research application.

3.6. Example of the Use of Regional Information in Nigeria

To justify our choice of providing two Human Lights products, we demonstrate how important the degree of regional information is in the filtering process. For this purpose, we give an example by running the process for Nigeria in the year 2001. The images in Figure 2a,b show the difference between including regional information or not. We also compare the classification to the Stable Lights and GRUMP data. In Figure 2a, we include the variables *avg_light* and *freq_light* as inputs to the random forest algorithm. In Figure 2b, we add a battery of regional variables (*lm_freq_5*, *lm_avg_25*, *lm_freq_49*, *lm_avg_99*, *local_detections_149* and *local_noise_199*). We set the total number of lit pixels to the same level as in the Stable Lights. Red and green colors highlight where the classification of Stable Lights and the Human Lights example specifications differ. Human Lights adds light in the case of green pixels and removes light in the case of red pixels. In the case of white pixels, Stable Lights and Human Lights are equal.

The main difference is in the southern part of Nigeria, which is affected by gas flaring. Stable Lights leave the light in a pixel if the average light and its frequency exceed a certain threshold. Therefore, it does not consider the differing signals at the top end of the light distributions. Consequently, all areas affected by gas flaring are lit. In contrast, our random forest classification algorithm recognizes patterns related to built-up areas at all levels of light distributions. Therefore, even the specification without the regional variables (Figure 2a) reduces the number of lit pixels in the gas flare region.

The light reduced mainly from the gas flare area is distributed very differently in the two examples. Note that we fixed the overall amount of light to the level of Stable Lights. In Figure 2a without regional variables, light is spread more evenly across the country. It distributes the light in pixels where the average light and frequency are similar to the built-up areas. In contrast, more lights in pixels close to cities are left on when using the regional variables (Figure 2b), even if the average light and frequency are relatively low. A visual inspection of GRUMP settlement points (yellow dots) hints that the regional information can help to improve the detection of population agglomerations, such as in the case of the cluster in the southwest.

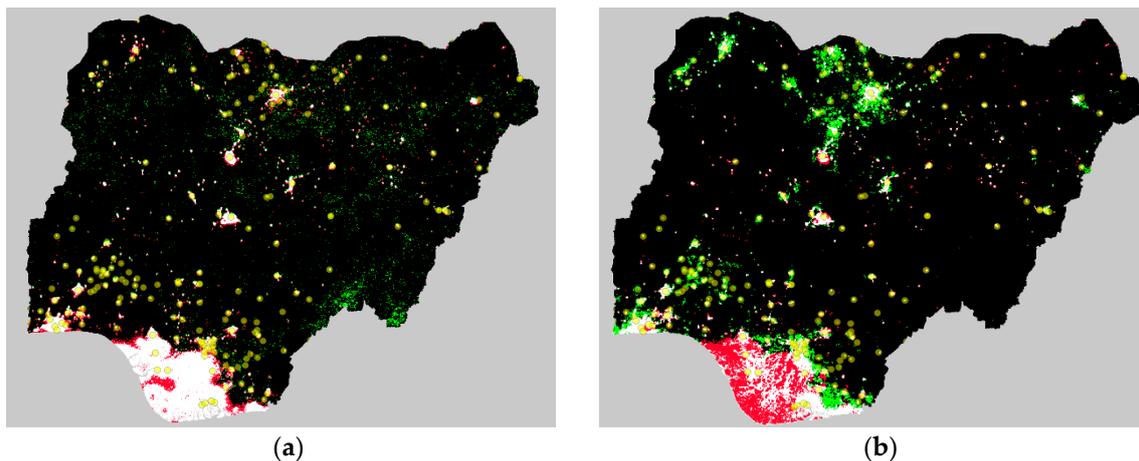


Figure 2. Impact of regional information on pixel classification in Nigeria. We use the satellite F15 in the year 2001. (a) no regional information (b) with regional information.

We also calculate statistics for the success in detecting built-up pixels and settlements in Nigeria for the Human Lights examples and Stable Lights. In Table 2, we summarize the results. Our results show that the filtering process without regional information performs worst (Column 2). We achieve the best results by including the regional variables (Column 3). Stable Lights, with some amount of regional information, lies between these. The results, together with the images, confirm the tradeoff between the accuracy in classification and the introduction of regional bias. If we utilize more regional information in the classification process, we treat pixels across regions in an increasingly different way, but at the same time, we gain an increasing advantage in human-generated light detection.

Table 2. Classification success in Nigeria.

	Stable Lights	Human Lights Figure 2a	Human Lights Figure 2b
	(1)	(2)	(3)
Pixels lit (%)	8.4	8.4	8.4
Built-up pixel detection (%)	75.1	73.5	80.2
Population settlement detection (%)	87.0	83.7	90.0

The Nigerian example shows that the choice of the degree of regional information has a direct impact on the results. Therefore, we provide two alternative approaches to classify the light in pixels as human-generated or background noise. The Global Human Lights product aims to reduce the regional bias, while the Local Human Lights product focuses on improved classification. Together with the Stable Lights, they can be used to evaluate the impact of regional information use on a given research application.

4. Results

We compare the Human Lights products with Stable Lights by visual inspection, the detection of built-up pixels, and the detection of settlement points from GRUMP. We provide the results for global and continental levels.

We first discuss the visual inspection of our results. Figure 3 shows the light distribution using the global approach; Figure 4 considers the local approach. In both cases, white pixels are classified as lit in both data sets; i.e., Stable Lights and Human Lights. The differences between the two data sets mark red and green pixels. Green pixels are lit in Human Lights, but dark in Stable Lights; red pixels are lit in Stable Lights, but dark in Human Lights. Please note that the images do not include any information about light intensity; we only compare in a binary sense whether Human Lights classifies a pixel as lit compared to Stable Lights or not.

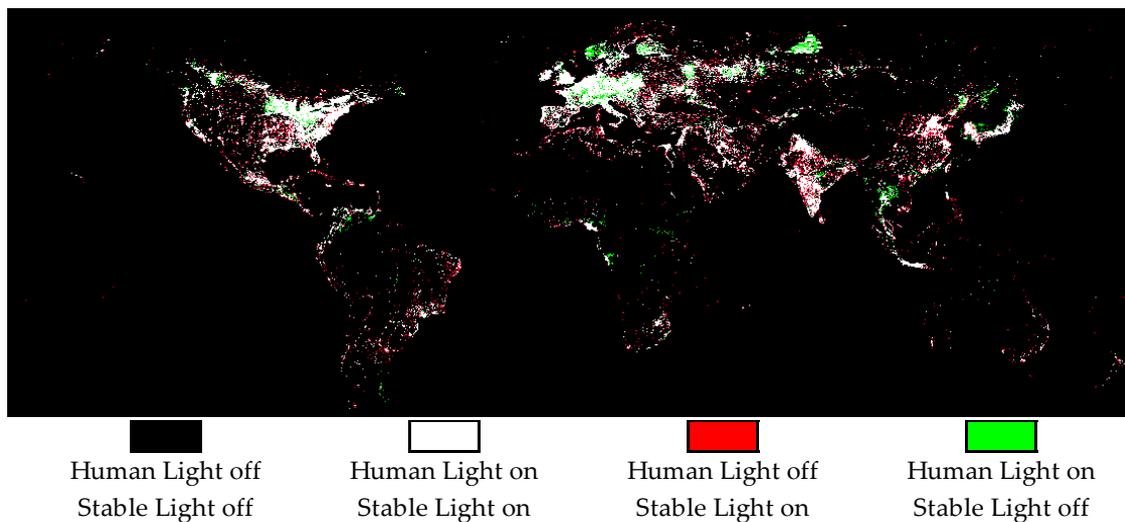


Figure 3. Global Human Lights vs. Stable Lights.

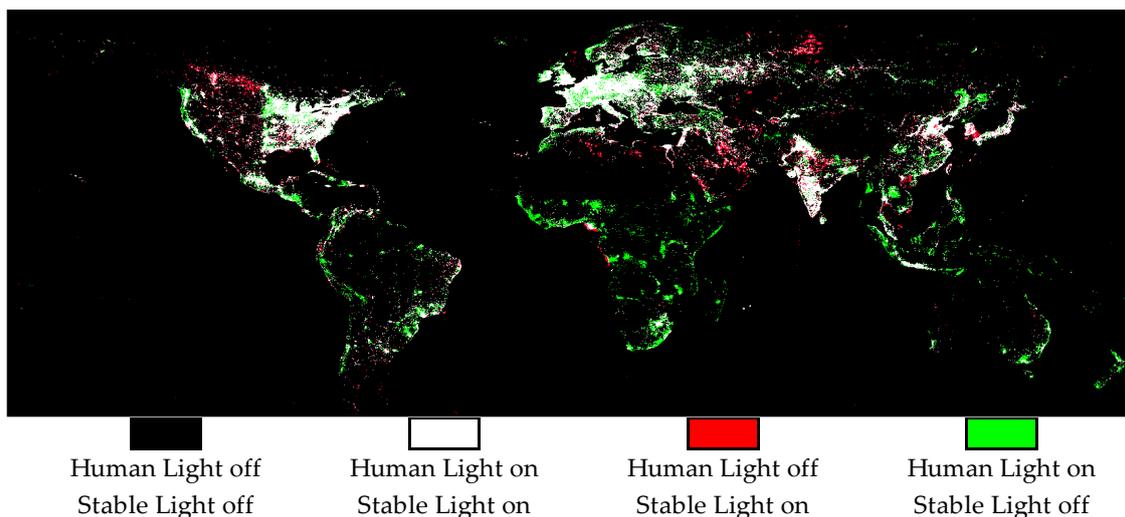


Figure 4. Local Human Lights vs. Stable Lights.

The first thing to note in the Global Human Lights product (Figure 3) is that it lights up fewer pixels than Stable Lights overall. The reason is that using the same filtering process for the whole world reduces the classification accuracy. This means that the ratio of misclassifications to correct classifications is relatively high if we add more pixels. As a result, we cannot add as many lit pixels as

those in the products that utilize regional information. In the Local Human Lights product (Figure 4), we are more confident about the classification accuracy, and therefore we can add more lit pixels overall.

The increase (Local Human Lights) or decrease (Global Human Lights) in lit pixels is not distributed evenly across regions. Therefore, we provide descriptive statistics at continental levels. The continent boundaries are based on the Natural Earth Admin 0 Countries dataset (Version 4.1.0 resolution 1:10 m). Table 3 reports the share of lit pixels for the three products. The statistics confirm the visual observation that the Global Human Lights has the lowest number of lit pixels and the Local Human Lights the highest. The performance of Stable Lights is between these.

Table 3. Number of lit pixels.

	Region	Stable Lights	Global Human Lights	Diff.	Local Human Lights	Diff.
		(1)	(2)	(3)	(4)	(5)
Pixels lit (%)	World	4.10	3.43	−0.67	4.77	0.68
	Africa	2.95	2.10	−0.85	8.25	5.30
	Asia	17.79	13.01	−4.78	18.98	1.19
	South America	6.22	4.17	−2.04	9.03	2.82
	North America	15.02	12.81	−2.21	15.26	0.24
	Europe	16.94	16.54	−0.39	19.12	2.18
	Oceania	2.10	1.01	−1.09	5.25	3.15

The statistics of the different continents provide interesting insights. Column 3 shows the difference in the share of lit pixels between Stable Lights and Global Human Lights. The share of lit pixels is reduced more for some continents than others. This finding indicates that the Stable Lights does not treat pixels with similar light characteristics in different regions equally. The logic is that the Global Human Lights product uses the same filtering process globally, and therefore similar pixels are treated more equally across regions. If the Stable Lights would also treat light in pixels independent of the region, then we would see relatively similar light reductions in all continents. Originally, the Stable Lights does not claim to treat regions in a uniform fashion, but this assumption is sometimes made in economic applications. We discuss this point later in more detail. The highest relative decrease is in Oceania, and the highest absolute decrease is in Asia. Our results suggest that these continents especially are favored in receiving a relatively high share of lit pixels in Stable Lights. Pixels in Europe, on the other hand, seem to require relatively high light values to be classified as stable light. However, we have to take this interpretation as a ballpark estimate because the two regional variables in the Global Human Lights product also lead to the differential treatment of pixels across regions.

Column 5 compares the difference between Stable Lights and Local Human Lights. The latter adds a relatively high number of lit pixels in Africa, South America and Oceania. This finding suggests that the Stable Light filtering process may be too strict, especially in these regions, when it comes to finding human-generated light. The low amount of added light in Asia and North America indicates that Stable Lights is quite successful in isolating human-generated light from background noise in these regions.

Table 3 evaluates the products in terms of correctly classifying built-up pixels. We remind the reader that the built-up data is not a perfect measure of success because it does not perfectly overlap with human-generated light; however, they allow us to give a rough comparison of how the different products detect human-generated light. Correct classifications refer to pixels being classified as lit when they have built-up content and not being lit when the built-up content is zero. The amount of correct classifications is divided by the total amount of pixels in a region to obtain the percentage of correct classifications. Overall, both Human Lights products improve the overall classification accuracy when compared to Stable Lights.

The classification accuracy depends on how successful a product is in leaving light on in built areas, but it is also strongly influenced by the number of lit pixels overall. Even if the lit pixels of a product would perfectly overlap with the built areas, the accuracy may be low if the number of lit pixels is very low. In contrast, if the number of lit pixels is too high relative to the number of built-up

pixels, then the accuracy is also reduced. In Table 4, we investigate the classification accuracy, without controlling for the amount of lit pixels. The classification accuracy is particularly higher in Asia for both Human Lights products. This is partly due to Stable Lights adding a relatively high share of lit pixels in the continent, but also partly attributable to erroneous classifications. Stable Lights outperforms both Human Lights products in Oceania. Adjusting the share of lit pixels regionally would lead to a better accuracy, but we refrain from doing so in order to keep the process reproducible.

Table 4. Correctly classified pixels.

		Region	Stable Lights	Global Human Lights	Diff.	Local Human Lights	Diff.
			(1)	(2)	(3)	(4)	(5)
Correct built-up pixelclassifications (%)	World		97.07	97.23	0.16	97.74	0.67
	Africa		92.87	92.91	0.04	94.48	1.61
	Asia		86.36	87.62	1.27	90.14	3.78
	South America		93.96	94.77	0.81	94.54	0.58
	North America		92.14	93.08	0.94	93.91	1.77
	Europe		91.21	90.71	−0.50	92.69	1.48
	Oceania		97.32	97.25	−0.06	97.14	−0.18

So far, we have focused on the classification accuracy in a binary sense; i.e., whether our data can predict areas with or without built-up content. Now, we add the information regarding the amount of built-up content in a pixel. For this purpose, we utilize the original GHSL data that record the built content in a pixel as a share of its total size. This allows us to investigate which part of the distribution of the built-up content the improvement in accuracy comes from. Figure 5 illustrates correct classification rates by different categories of built-up content for the whole world (Figure 5a) and Africa (Figure 5b), where the Local Human Lights product performs remarkably well. The share of lit pixels in Africa increases from 2.1% in Stable Lights to 8.25% in Local Human Lights. Despite the huge increase, the classification accuracy is 1.61% better for Local Human Lights. This indicates that the added light is not simply randomly distributed but actually lands in pixels with built-up content. We achieve the largest gain in accuracy in areas that do not have high built-up content. This is a very important result because our aim is to improve human-generated light detection especially in marginal areas where it is difficult to distinguish human-generated light from background noise.

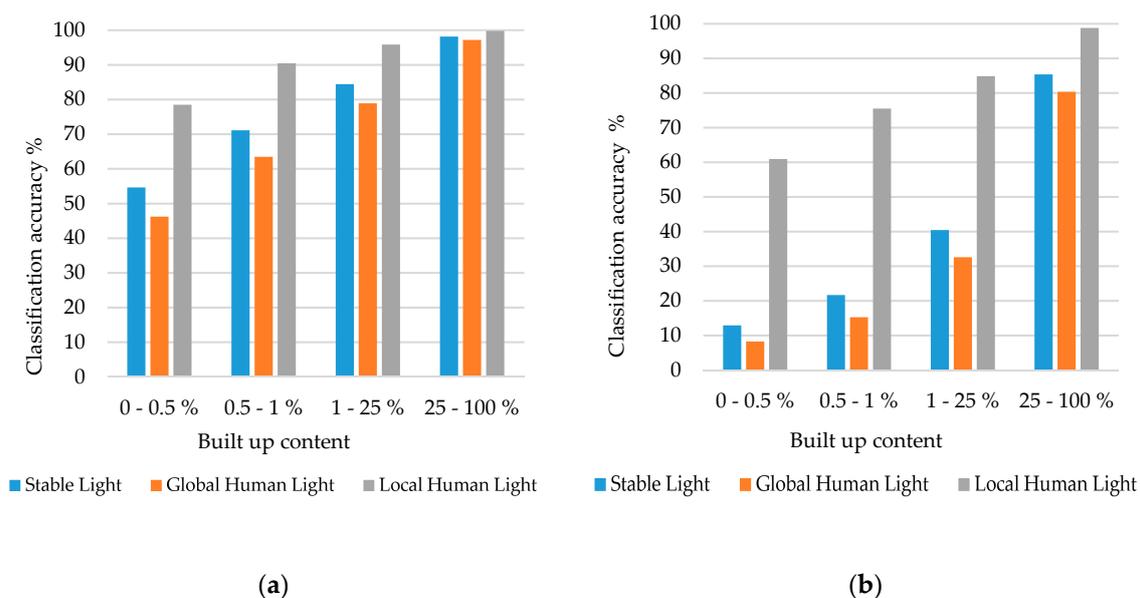


Figure 5. Correctly classified pixels by built-up content. (a) World; (b) Africa.

Finally, we analyze which data has the highest success rate in detecting settlement points in the GRUMP data in Table 5. Note that this is an external data source that we have not used in the

data-generating process. Global Human Lights performs worse than Stable Lights. This is mainly due to the lower share of lit pixels in the Global Human Lights product, but is also related to Stable Lights being very successful in finding the settlements. A visual inspection shows that lit pixels in Stable Lights coincide with population settlement points, with the exception of gas flare areas. However, making a minimal number of mistakes suggests that it might be possible to add lit pixels with a high success rate. This is done by the Local Human Lights product in Column (4) and (5). Overall, Local Human Lights improves the detection rate of population settlements by 1.05%. The most notable gain is in Africa, where the detection rate increases by 15.3%. This finding supports our argument that the Stable Lights filtering process is too strict and therefore misses light from economic activity, especially in developing countries.

Table 5. Population settlement detections.

Region		Stable Lights	Global Human Lights	Diff.	Local Human Lights	Diff.
		(1)	(2)	(3)	(4)	(5)
Population settlements detected (%)	World	93.29	89.32	−3.97	94.34	1.05
	Africa	66.32	59.39	−6.93	81.62	15.30
	Asia	87.75	79.75	−8.01	88.00	0.25
	South America	91.78	87.05	−4.74	94.39	2.61
	North America	98.19	95.95	−2.23	98.13	−0.06
	Europe	99.11	98.60	−0.51	99.25	0.14
	Oceania	95.48	91.22	−4.26	96.81	1.33

The results so far are mostly driven by the differences in the number of lit pixels in the different products. Table 6 sets the number of pixels in the Human Lights products as equal to Stable Lights by continent to evaluate their performance. Remember that we do not advise doing so in data applications because we do not know exactly how the amount of light is set in the Stable Lights. Here, we aim to identify whether the improvement of detection rates comes from the amount of light or its distribution. We first concentrate on the GHSL data (Columns 1–3). Column 2 shows that Global Human Lights has a worse built-up classification accuracy in all continents. In contrast, Local Human Lights outperforms Stable Lights in every continent, as shown by the results reported in Column 3. Interestingly, the order of performance of the different products reflects the degree of regional information in each product. Given the same amount of light, increasing the use of regional information leads to improved built-up classification accuracy.

Table 6. Performance analysis; light amount fixed to Stable Lights.

Region	Correct Built-up Pixel Classification (%)			Population Settlements Detection (%)		
	Stable Lights	Global Human Lights	Local Human Lights	Stable Lights	Global Human Lights	Local Human Lights
	(1)	(2)	(3)	(4)	(5)	(6)
Africa	92.87	92.58	94.53	66.32	62.89	65.79
Asia	86.36	86.22	90.34	87.75	86.60	86.86
South America	93.96	93.40	95.39	91.78	89.71	91.88
North America	92.14	91.97	93.99	98.19	96.96	98.06
Europe	91.21	90.57	93.49	99.11	98.65	99.11
Oceania	97.32	97.05	98.04	95.48	93.88	89.36

Next, we evaluate the data based on GRUMP (Columns 4–6). We compare the population settlement detection rates while fixing the amount of light to the Stable Lights level for each continent. Here, we see that Stable Lights outperforms Human Lights in most cases. However, Local Human Lights compares very well considering that it lacks the data on daily light values and subjective regional thresholds that are used in Stable Lights. Together with the findings above, we conclude that Stable Lights performs very well in finding light from economic activity; it is simply missing a considerable part of it.

The results reported in Table 6 also suggest that the performance of the products depends on the data that we use as a yardstick. On the one hand, built-up data have the disadvantage that we use them to train our random forest algorithm and that they are not a perfect label for light generated in economic activity. However, they do have a global coverage and also include the amount of light that is erroneously classified as human-generated. On the other hand, population settlement data are not used in the Human Lights process, but they do not provide us information regarding how the products perform at the margin; they only cover specific points and do not show the downside of increasing the amount of lit pixels.

5. Discussion

We provide the Human Lights products as easily accessible benchmarks. However, as we discuss in Section 3, the products contain many impactful filtering process settings. In particular, the amount of regional information, the number of lit pixels and the parameters in the random forest algorithm are settings that can be set more optimally to a specific research application. Therefore, we discuss the implications of the different approaches in common economic research applications.

The most important economic research applications use nightlights as a proxy for GDP. However, comparing aggregates on the country level assumes that the measurement errors of the two indicators are uncorrelated. If the measurement errors of the two indicators are uncorrelated, then they can complement each other to give a more precise estimate of true economic activity. Existing studies acknowledge that the measurement error in GDP data is associated with the level of GDP itself given differences in the capacity of statistical authorities [1,5]. For nightlight data, they make the seemingly reasonable assumption that the errors are uncorrelated with the errors in GDP, because the light images are taken from space in a uniform fashion for the whole world. Even if there were some errors in the imaging process, they show up in the satellite light image independent of GDP. We agree that the assumption holds for the raw satellite images. However, as we argue in Section 3, this is no longer true after using a regionally varying tool for the filtering of background noise. Our local product and the Stable Lights are fully exposed to this effect, but we try to mitigate it with the global product. We suggest using the different products together in order to provide a more complete picture of the connection between lights and GDP.

The insight that a filtering process does not classify pixels with the same light values uniformly across the globe can also be addressed with econometric methods. The easiest solution is to focus on changes in light and GDP instead of levels [5]. This approach eliminates systematic time-invariant differences in classification across regions. However, the time dimension adds a challenge because the sensor light gain setting in the satellite imaging process changes over time [24]. As a result, overall light intensity varies strongly across years. The current best practice in economics is to use satellite and/or year-fixed effects in regression analysis to control for the changes in the overall light intensity. The assumption is that all regions are affected equally by the overall increase in light. We agree again that this may be reasonable for the raw images, but we do not know how strongly this holds after applying local filter methods. Our argument is that a locally varying threshold to set the light amount makes sure that an increase in overall light is unevenly distributed across regions. As an example, imagine that we would add 1 to the light value of every pixel in the average visible band image; two locations with identical light characteristics in different regions can be pushed above the filtering threshold because of the overall increase in light. However, if the filtering threshold differs regionally, then only one of the locations might be classified as lit.

Economic applications that refrain from drawing a global functional link between light and a variable of interest have an additional tool to address the regional differences in light. They can use country or region-specific fixed effects in addition to the satellite/year-fixed effects to control for time and space-invariant heterogeneities. The assumption is that, in a chosen region, the overall light increase has the same effect across years. This assumption is only valid if the filtering process settings are fixed in the region over time. In our local product, the settings are exactly the same for

the sub-regions of 250×250 pixels. However, in Stable Lights, the sub-regions—where all pixels are affected by the same light-free thresholds—are much smaller, with a size only just 25×25 pixels. If a region under study is larger than those areas, then light increases might be distributed unevenly. Therefore, we encourage researchers to run the filtering process based on the regional units in a study (e.g., countries, administrative regions or grid cells). Such an approach, together with region and year-fixed effects in a regression analysis, will maximize the accuracy in human-generated light detection and minimize regional biases.

Despite the challenges in filtering out background noise, light data are still an invaluable indicator of regional development. Our aim is to draw attention to the possible biases that should be addressed in the context of a given research application. We provide the benchmark ready-to-use Human Lights products that facilitate robustness analysis. However, the best results for separating human-generated light from background noise and understanding the magnitude of the bias can only be achieved by tailoring the process to fit the needs of the research application. The research applications vary in their demand for classification precision and the size of the study region, which is why it is not possible to provide a dataset that fits all needs.

6. Conclusions

In this paper, we provide an alternative approach to filter out background noise from satellite nightlight images. Currently, the most widely used nightlight product is Stable Lights, which filters out transient light sources from the average visible band image. Stable Lights has been widely adopted in economics as a proxy for economic activity because researchers are free to choose their region of interest, and the data quality is constant across regions. Our approach differs in that we try to separate light sources based on whether they are of human origin or not. We argue that human-generated light, even if it comes from a transient light source, is a better proxy for economic activity than only relying on stable light sources. In particular, rural areas or small cities in developing countries might have only temporary access to the electric grid, reducing the probability of being identified in Stable Lights. Our alternative approach helps us to improve the detection of economic activity, especially in the lower tail of the light distribution. We aim to improve the quality of nightlight data, especially in developing countries where other data sources lack quality.

We apply the random forest machine learning algorithm to classify light in pixels as human-generated or not. It learns patterns in the light data that are associated with human-generated light, which we proxy using built-up areas. The algorithm then uses the learned patterns to predict the probability that light in a pixel is of human origin. Finally, we use the predicted probabilities to add the light in a way that most likely chooses human-generated light over background noise.

A background noise filtering process contains inevitable choices on the degree of regional information and the number of lit pixels. The optimality of the choices depends on a given data application because they trade off data suitability to improvements in human-generated light detection. These choices are arbitrary if the resulting product is designed for general use. We show how these choices have a direct impact on the results when applying the data. In economic applications, the issues of correlated measurement error and spatial correlation should be addressed in particular. We provide two different benchmark products based on these choices. The Human Lights products can be used together with the Stable Lights to evaluate the impact on a given data application. We also provide instructions in the supplementary materials regarding how to redo or adjust the Human Lights filtering process, which can be used to optimize the choice of regional information and the amount of light for a given data application.

Our main findings are that the Stable Lights product does not detect considerable amounts of light from economic activity at the lower tail of the light distribution, and that the missing light from economic activity is not equally distributed across regions. Using the Human Lights products, we demonstrate that the missing economic activity is mostly located in developing countries, where the nightlight data is most useful due to lacking official statistical data of sufficient quality. Furthermore,

we show that it is important to be transparent regarding the process of filtering background noise and to evaluate this process since it directly affects the results of research applications. Both Human and Stable Lights contain arbitrary choices, such as the choice of the classification rule, regional light variables and regional window sizes. In Section 3, we provide reasoning for our choices, but we note that they might also not be optimal in all data applications. Therefore, we invite discussion on the inevitable arbitrary choices in the filtering process to improve the best practice.

The Human Lights products presented in this paper are available at GESIS datorium. In our future work, we aim to expand the products for all satellite years and different regional units (e.g., countries). For now, researchers interested in applying the Human Lights method can follow the R code in the supplementary materials. Globally, the code takes weeks to run for only one satellite/year on a workstation (32 cores, 128 GB RAM) even if parallel computing is utilized. Therefore, it is cumbersome to adjust the code and to run it on a global level for all years. For smaller regions, such as a single country, the process runs much more quickly.

Recently, new and more accurate remotely sensed data is available to approximate economic activity. For example, the Visible Infrared Imaging Radiometer Suite (VIIRS) from NOAA provides new, improved nightlight images [25]. In addition, daylight images facilitate a more accurate measurement of economic activity [26]. However, the new data sources are only available for recent years. The value of DMPS-OLS nightlight data in economics is the long time-span that allows the study of longer-term development. Therefore, we encourage its further development as a measure of economic activity.

Supplementary Materials: The code and data are permanently published on GESIS datorium, DOI: dx.doi.org/10.7802/1918.

Author Contributions: Conceptualization, I.M. and C.L.; methodology, I.M.; software, I.M.; validation, I.M.; formal analysis, I.M.; investigation, I.M.; resources, I.M. and C.L.; data curation, I.M.; writing—original draft preparation, I.M. and C.L.; writing—review and editing, I.M. and C.L.; visualization, I.M.; supervision, C.L.; project administration, C.L.; funding acquisition, C.L.

Funding: This project is funded by the German Research Foundation (DFG) grant No. BL 1502/1-1. We acknowledge support by the Open Access Publication Funds of the Technische Universität Braunschweig.

Acknowledgments: The authors are grateful for many helpful comments by Richard Bluhm, Melanie Krause, and the participants of the Workshop on Geodata and Economics at TU Braunschweig in 2018.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

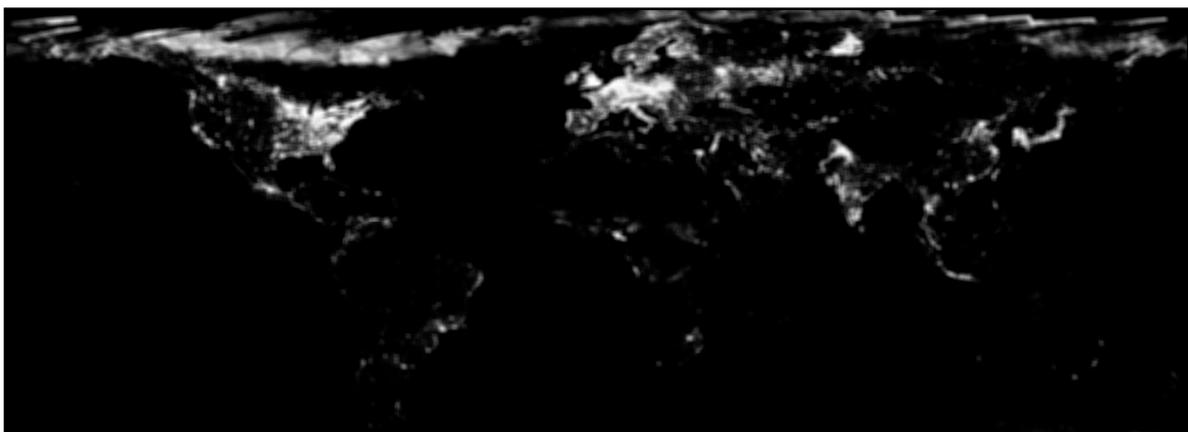


Figure A1. Local noise

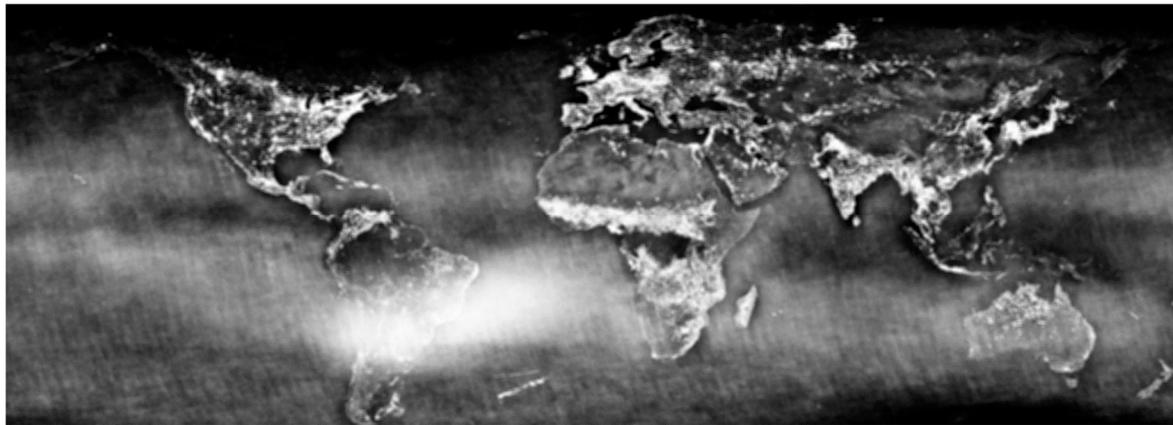


Figure A2. Local light detection

References

- Chen, X.; Nordhaus, W.D. Using luminosity data as a proxy for economic statistics. *PNAS* **2011**, *108*, 8589–8594. [[CrossRef](#)] [[PubMed](#)]
- Sutton, P.C.; Costanza, R. Global estimates of market and non-market values derived from nighttime satellite imagery, land cover, and ecosystem service valuation. *Ecol. Econ.* **2002**, *41*, 509–527. [[CrossRef](#)]
- Sutton, P.C.; Elvidge, C.; Ghosh, T. Estimation of gross domestic product at sub-national scales using nighttime satellite imagery. *Int. J. Ecol. Econ. Stat.* **2007**, *8*, 5–21.
- Elvidge, C.; Sutton, P.C.; Ghosh, T.; Tuttle, B.T.; Baugh, K.E.; Bhaduri, B. A global poverty map derived from satellite data. *Comput. Geosci.* **2009**, *35*, 1652–1660. [[CrossRef](#)]
- Henderson, J.V.; Storeygard, A.; Weil, D.N. Measuring economic growth from outer space. *Am. Econ. Rev.* **2012**, *102*, 994–1028. [[CrossRef](#)] [[PubMed](#)]
- Henderson, J.V.; Squires, T.; Storeygard, A.; Weil, D. The global distribution of economic activity: Nature, history, and the role of trade. *Q. J. Econ.* **2017**, *133*, 357–406. [[CrossRef](#)]
- Storeygard, A. Farther on down the road: Transport costs, trade and urban growth in sub-Saharan Africa. *Rev. Econ. Stud.* **2016**, *83*, 1263–1295. [[CrossRef](#)] [[PubMed](#)]
- Michalopoulos, S.; Papaioannou, E. National institutions and subnational development in Africa. *Q. J. Econ.* **2013**, *129*, 151–213. [[CrossRef](#)] [[PubMed](#)]
- Alesina, A.; Michalopoulos, S.; Papaioannou, E. Ethnic Inequality. *J. Political Econ.* **2016**, *124*, 428–488. [[CrossRef](#)]
- Lessmann, C.; Seidel, A. Regional inequality, convergence, and its determinants—A view from outer space. *Eur. Econ. Rev.* **2017**, *92*, 110–132. [[CrossRef](#)]
- Hodler, R.; Raschky, P. Regional Favoritism. *Q. J. Econ.* **2014**, *129*, 994–1033. [[CrossRef](#)]
- Lessmann, C.; Steinkraus, A. The Geography of Natural Resources, Ethnic Inequality and Civil Conflict. *Eur. J. Polit. Econ.* **2019**, *59*, 33–51. [[CrossRef](#)]
- Fabian, M.; Lessmann, C.; Sofke, T. Natural Disasters and Regional development—The Case of Earthquakes. *Environ. Dev. Econ.* **2019**, *24*, 479–505. [[CrossRef](#)]
- Baugh, K.; Elvidge, C.; Ghosh, T.; Ziskin, D. Development of a 2009 stable lights product using DMPS-OLS data. *Proc. Asia Pac. Adv. Netw.* **2010**, *30*, 114–130. [[CrossRef](#)]
- Bluhm, R.; Krause, M. *Top Lights—Bright Cities and Their Contribution to Economic Development*; CESifo Working Paper Series No 7411; Center for Economic Studies: Munich, Germany, 2018.
- Abrahams, A.; Oram, C.; Lozano-Gracia, N. Deblurring DMSP nighttime lights: A new method using Gaussian filters and frequencies of illumination. *Remote Sens. Environ.* **2018**, *210*, 242–258. [[CrossRef](#)]
- Bruederle, A.; Hodler, R. Nighttime lights as a proxy for human development at the local level. *PLoS ONE* **2018**, *13*, e0202231. [[CrossRef](#)] [[PubMed](#)]
- Pesaresi, M.; Ehrlich, D.; Ferri, S.; Florczyk, A.J.; Freire, S.; Halkia, S.; Julea, A.M.; Kemper, T.; Soille, P.; Syrris, V. *Operating Procedure for the Production of the Global Human Settlement Layer from Landsat Data of the Epochs 1975, 1990, 2000 and 2014*; Publications Office of the European Union: Brussels, Belgium, 2016. [[CrossRef](#)]

19. Elvidge, C.D.; Baugh, K.E.; Kihn, E.A.; Kroehl, H.W.; Davis, E.R. Mapping city lights with nighttime data from the DMSP operational linescan system. *Photogramm. Eng. Remote Sens.* **1997**, *63*, 727–734.
20. Croft, T.A. Burning waste gas in oil fields. *Nature* **1973**, *245*, 375–376. [[CrossRef](#)]
21. Ghosh, T.; Anderson, S.; Powell, R.; Sutton, P.; Elvidge, C. Estimation of Mexico’s informal economy and remittances using nighttime imagery. *Remote Sens.* **2009**, *1*, 418–444. [[CrossRef](#)]
22. Martinez, L.R. *How Much Should We Trust the Dictator’s GDP Growth Estimates?* Working Paper; Harris School of Public Policy, University of Chicago; Social Science Research Network: Rochester, NY, USA, 2019. [[CrossRef](#)]
23. James, G.; Witten, D.; Hastie, T.; Tibshiriani, R. *An Introduction to Statistical Learning*; Springer: New York, NY, USA, 2013; ISBN 1461471370.
24. Elvidge, C.; Hsu, F.C.; Baugh, K.; Ghosh, T. National trends in satellite-observed lighting. *Glob. Urban Monit. Assess. Earth Obs.* **2014**, *23*, 97–118.
25. Chen, X.; Nordhaus, W.D. VIIRS nighttime lights in the estimation of cross-sectional and time-series GDP. *Remote Sens.* **2019**, *11*, 1057. [[CrossRef](#)]
26. Jean, N.; Burke, M.; Xie, M.; Matthew, W.; Davis, D.B.; Lobell, B.; Ermon, S. Combining satellite imagery and machine learning to predict poverty. *Science* **2016**, *353*, 790–794. [[CrossRef](#)] [[PubMed](#)]



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).