Title: Biodiversity Conservation Can Coexist with Increasing Human Footprint

Authors: Patrick W. Keys [1, *], Elizabeth A. Barnes [2], Neil Carter [3]

Affiliations:

[1]School of Global Environmental Sustainability, Colorado State University, Fort Collins, CO

[2]Department of Atmospheric Science, Colorado State University, Fort Collins, CO

[3]School for Environment and Sustainability, University of Michigan, Ann Arbor, MI

* corresponding author: Patrick Keys, patrick.keys@colostate.edu

Author contributions:

All authors devised the research plan. E.A.B. undertook the development of the convolutional neural network, including design, training and validation. P.W.K. processed sustainable development goal data and performed quality control on neural network model output. All authors contributed to the writing of the manuscript.

Abstract

Fundamental to the success of sustainable development is a foundation of intact ecosystems. While the United Nations Sustainable Development Goal 15, "Life on Land", seeks to protect biodiversity in terrestrial ecosystems, accelerating human-driven changes across the Earth system are undermining efforts to preserve biodiversity. Understanding this tension has never been more urgent and requires tools that reveal pathways for development that also support biodiversity. Here we introduce a near-present, global-scale machine learning-based human footprint index which is capable of routine update. By comparing global changes in the machine learning human footprint index between 2000 and 2019 to national-scale biodiversity metrics for Goal 15, we find that some countries are experiencing increases in their human footprint while biodiversity metrics are improving as well. We further examine development and policy dynamics to uncover enabling mechanisms for balancing increased human pressure with biodiversity gains. This has immediate policy relevance, since the majority of countries globally are not on track to achieve Goal 15 by the declared deadline of 2030. Moving forward, the machine learning human footprint index can be used for ongoing monitoring and evaluation support toward the twin goals of fostering a thriving society and global Earth system.

Introduction

The human footprint index (HFI) represents one of the most important tools for interpreting human pressure on the landscape(1). A dimensionless metric which captures the extent of human influence on the terrestrial surface, the HFI is distinct from many land-use metrics in that it captures the total influence of human existence on a given location in a single index, rather than categorizing individual land use types (2). The applications to which the HFI is put to use are manifold (3–6), with enormous demand for information about human pressure on the land surface to support policies related to land use change, biodiversity conservation, and climate action (6–8). A key challenge for expanded use of the HFI for operational efforts, however, is that new updates typically lag the present day by seven or more years (9, 10). This time lag means that large-scale or substantial changes to the land surface caused by human activities can occur well before we detect them, hampering our ability to monitor and respond to their effects

on biodiversity. This temporal mismatch pertains directly to the achievement of the United Nations Sustainable Development Goals (SDG) since the deadline for their achievement is 2030 (11–13), with some specific targets within SDG15, "Life on Land", set for even sooner (14). While international (14) and voluntary national-scale assessments of various biodiversity metrics (15) are available for monitoring progress toward SDG15, comparing those SDG15 assessments with an operational interpretation of human influence on the terrestrial land surface - one that can be easily and regularly updated - would provide deeper insights on the social-ecological outcomes and processes of landscape change.

Machine learning, specifically convolutional neural networks (CNN), are well suited to the task of identifying patterns in imagery(16), including patterns of human activity across the land surface (17–19). We present the first near-present, global, machine learning-based HFI (ml-HFI) for the years 2000 and 2019 (Fig 1). We train a CNN to ingest space-borne imagery of the earth's surface and predict the most up-to-date version of the HFI for the year 2000 (Fig 1a; hereafter the Williams-HFI) (10). The conventional approach for creating the HFI requires a harmonization of eight different sub indices representing different aspects of human pressures on the terrestrial surface of the Earth, including built infrastructure, population density, and a variety of land use and land cover data. This resulting HFI is a dimensionless index ranging from a low to high human footprint. The Hansen Global Forest Change imagery version 1.7 (GFCv1.7)(20) includes global, cloud-free, growing season Landsat imagery for the years 2000 and 2019 (20), and so we exploit the overlap between the year-2000 Williams-HFI and the year-2000 GFCv1.7 imagery for training. We then use the trained CNN to predict the 2019 ml-HFI based on 2019 GFCv1.7 imagery.

Results

Human pressure increasing globally

Global, regional, and local patterns of human footprint from the Williams-HFI (Fig 1a) are identified in the year 2000 ml-HFI with high fidelity (Fig 1b). High accuracy is realized over all levels of human footprint by training on only 3% of the terrestrial locations (Supp Fig 3), with a

mean absolute error on unseen testing data of 0.07 (Supp Fig 6). With that said, there are regions of the planet that are systematically more challenging for the CNN to reproduce (Supp Fig 8), including rocky areas of remote deserts and high latitude mountain ranges. Training the model on a larger portion of locations is expected to overcome this limitation in future iterations. Once trained, the CNN is then tasked with predicting the most current version of the ml-HFI from 2019 GFCv1.7 imagery (Fig 1c). This 2019 map represents the most current quantification of the human footprint index, previously estimated for 2013 (10).

Changes in ml-HFI between 2000 and 2019 are visualized globally in Fig 2 with positive values indicating a large increase in HFI, and negative values indicating large decreases in the HFI (where "large" is defined as changes greater than 0.25). Patterns of change are consistent with a steady expansion of human encroachment into low HFI areas, as well as increasing density of pressure in areas of higher HFI. Inset panels in Fig 2 display the GFCv1.7 imagery from 2000 (left panels) and 2019 (middle panels) for different forms of increasing human pressure. While the human eye can easily see the changes in the images from one year to the next, we leverage a neural network visualization technique called Layerwise Relevance Propagation (LRP)(21, 22) to reveal the features that the neural network deems most relevant for its calculation of the ml-HFI (right panels). Inspection of the LRP maps indicates that the CNN is correctly identifying human-made features and activities (e.g. transportation networks or land-cover conversion) providing confidence in, and interpretability of, the CNN's predictions. This step is especially important if the ml-HFI is to be used for policy decisions as national and international laws have begun to require explainable AI systems for decision-making (23, 24).

Biodiversity gains despite development

Previous research underscores the transformative effects that human activities have on natural ecosystems (25, 26). Many of these studies indicate that heightened human pressure on the landscape diminishes the capacity for those lands to support biodiversity. We interrogate those assumptions using the ml-HFI for the years 2000 and 2019 with a contemporary assessment of country-level progress toward achievement of SDG15 (14). We find that 76 countries are

experiencing substantial increases in human pressure together with stagnating or decreasing progress toward SDG15. This supports the common belief that an increasing human footprint is not compatible with the maintenance or improvement of terrestrial biodiversity. Contrary to this common belief, however, we also find that 43 countries (Fig 2, green outlines) are experiencing substantial increases in human pressure over a large fraction of their land surface (increases of 0.25 or larger over 1% or more of their area) while at the same time "moderately improving" or actively "on track to achieve" SDG15. These substantial changes identified by the ml-HFI have been further confirmed manually for more than 2,000 locations across the 43 countries using the time lapse feature in Google Earth Pro.

Our core finding, that the ml-HFI is increasing while biodiversity gains are also increasing, remains nascent in the conservation and development discourse (4, 27–29). To better understand this coexistence, we summarize the types of changes detected with the ml-HFI along with specific SDG15 indicator progress in Fig 3. The ml-HFI reveals a wide spectrum of human pressure on the landscape, including: deforestation for extensive mining (Suriname, Guyana), spreading road networks (Estonia, Slovenia), urbanization (Uganda, The Gambia), and expanding agriculture (Benin, Morocco). Moreover, there is no uniform pathway for making progress across the SDG15 indicators for forest-(30), freshwater-(31, 32), and threatened "red list" species conservation (33). For example, Slovenia and Estonia are on track to achieve SDG15 in all three indicators. But for most countries, it is a very mixed pathway across the three indicators. For example, Uganda may be making progress on forest and freshwater conservation, while decreasing toward its "red list" species target. Evidently, the types of increasing pressures revealed by the ml-HFI do not lead to identical patterns of SDG15 progress. This suggests additional country-level policy information may deepen our understanding, and we focus on three cases: Guyana, The Gambia, and Morocco.

Discussion

Cases of development processes

A core function of the ml-HFI is to monitor remote regions experiencing rapid change, as clearly

identified in Guyana - including considerable expansion of road networks, deforestation and extensive mining in the forested interior (Fig 3, Supp. Fig. 15). Notwithstanding its accelerating economic development, Guyana has the second-highest per capita forest cover in the world, second only to its neighbor, Suriname (15). Given the enormous potential for forest preservation, Guyana was one of the first countries to implement REDD+ (Reduced Emissions from Deforestation and Forest Degradation), partnering with Norway which provides financial resources to incentivize the reduction in deforestation rates, and the establishment of protected areas in collaboration with indigenous communities (15). Yet, the growing network of roads and mining across the Guiana Shield revealed by the ml-HFI shows how quickly development processes can encroach into wild and remote areas.

Urbanization of agricultural areas is a key feature of development (34), and as evidenced by the ml-HFI results for The Gambia, can take place explosively over the course of less than 20 years (Fig 3, Supp. Fig. 16). Elsewhere in The Gambia, the ml-HFI reveals changes driven by deforestation for roads and agriculture. The Gambia's Voluntary National Review for SDG15 highlights the role that climate change has on exacerbating challenges to natural resource management, as well as the tension between preserving biodiversity with food security for its population - which remains 50% rural (35). Nonetheless, The Gambia includes multiple protected forests, including biodiverse areas near the mouth of the The Gambia River. The ml-HFI results for 2019 show the lowest levels of human pressure in precisely these riparian locations, emphasizing the need for continued protection of these forests.

Inspection of the ml-HFI shows that Morocco is experiencing development pressure in the form of expanding peri-urban farmland, as well as road and energy infrastructure (Fig 3, Supp. Fig. 17). As a middle income country, with both semi-arid forests and desert ecosystems, Morocco has focused on boundary demarcation of forested areas, enforcement of forest protection from illegal activities, and expanding coverage of forest management plans (36). As with Guyana and The Gambia, development is still taking place and, as clarified in Morocco's Voluntary National Review for the SDGs, there are ongoing plans for water and land development as well as forest

conservation in Morocco (36). The ml-HFI results for 2019 support the assessment that considerable potential exists for conservation, especially in the remote mountains and arid ecosystems of the east and southern reaches of Morocco.

Prospects for development monitoring

Combining the SDG15 progress with our 2019 present-day human footprint index shows that progress toward biodiversity goals can be made despite increases in human pressure. These increases occur in countries ranging in wealth, economic development, or reliance on specific natural resources. For many countries that are aiming for resource intensive economic development, monitoring and evaluation will be critical for tracking changes in near-real time. Countries that are still actively developing may find the ml-HFI useful to identify strategies to provide effective guardrails for staying on track for protecting biodiversity (e.g. Dominican Republic, Republic of Congo), especially when increased human activity encroaches on pockets of biodiversity. For wealthier or more developed countries with longstanding regulatory structures for protected areas, the ml-HFI can help determine conservation policy efficacy and ensure that additional increases in HFI concentrate in already developed areas (e.g. Sweden, Switzerland).

Although we detail compelling evidence of progress toward biodiversity in the face of increasing human pressure for some countries, the more common pattern is that of increasing human pressure at the expense of gains in biodiversity conservation. We do not single out any specific countries that are not on track to achieve SDG15 (14), but the changes in their ml-HFI results reveal similar pressures to the cases detailed above - urban growth, agricultural expansion, and resource extraction (mining and timber harvesting). While considerable efforts are being made to protect intact ecosystems from encroaching or intensifying human activities, growth in human influence is likely outpacing field-based monitoring protocols to assess environmental impacts and highlight that many current policies are insufficient to regulate development (30). Capable of capturing near present changes to the land surface, the ml-HFI could be an important tool for assisting countries that are flagging behind their SDG15 targets to close the gap by 2030.

That some countries can simultaneously be progressing toward biodiversity goals while experiencing dramatic increases in human pressure confounds expectations. Moving forward, science and policy must work to understand why and how this is possible. This will permit the development of better policies that will foster a thriving society and environment in the Anthropocene.

Methods

Data

Human Footprint Index

We quantify human pressure on the landscape using the human footprint index (HFI), which is a global, dimensionless index of human pressure on the global land surface (9, 37). We employ the most up-to-date version of the HFI (10), which we denote as "Williams-HFI" throughout this text. This version of the HFI is comprised of eight different sub indices representing different aspects of human pressures to the terrestrial surface of the Earth, including 1) extent of the built environment, 2) population density, 3) electric infrastructure, 4) agricultural lands, 5) pasture lands, 6) roadways, 7) railways, and 8) navigable waterways. We rescale the original index, which ranges from 0 to 50, to a range from 0 to 1.

Landsat imagery

We use the Global Forest Change dataset, which is a global analysis of forest cover change based on Landsat imagery (20). The data consist of processed Landsat imagery from the bands 3, 4, 5, and 7. The global processing identifies growing-season imagery and only includes cloud-free images. We specifically use bands 3, 5, and 7 in our analysis (corresponding to the "red" band, and two bands of "near infrared"). We chose these bands based on preliminary sensitivity tests, which revealed these were the most sensitive to interpreting the terrestrial changes we needed to identify. In all images shown here, the three GFCv1.7 channels are combined to create a false color image as near infrared bands are outside of the visible spectrum.

The latest version of the Global Forest Change product includes cloud-free processed Landsat

imagery for the years 2000 and 2019 (GFCv1.7). However, when cloud-free imagery was not available for the specific year, imagery was taken from the closest year with cloud-free data. For our 2000 analysis, this means imagery could come from years within the range of 1999-2012. For our 2019 analysis, this means imagery could come from years within the range of 2010-2015.

Convolutional Neural Network

Architecture

We train a 2D convolutional neural network (CNN) with 5 convolutional layers with zero padding and stride of 2, and a final hidden dense layer of 32 units that feeds into the output layer of size 1 (a single predicted HFI value). The first 4 convolutional layers are each followed by max pooling layers with pooling size 2x2, stride of 2, and valid padding. A dropout layer with a 50% dropout rate is applied following the final convolutional layer, prior to the dense layer. The first convolutional layer has 64 filters, while all subsequent convolutional layers have 128. All filters have size 3x3. All activations use the rectified linear unit (ReLu) except for the dense layer, which utilizes a sigmoid activation function to constrain the output between 0 and 1.

The input into the CNN is an array with shape (120,120,3) representing 3 Landsat channels of 120 pixels x 120 pixels each. An HFI pixel is approximately 40x40 Landsat pixels, however, we found that the model performed best when the input included surrounding Landsat pixels, likely to give the CNN more "context" to predict the HFI in the center of the image. Thus, we input 3 120x120 pixel Landsat images and task the CNN to output a single value between 0 and 1. This value represents the HFI in the center (see white boxes in Supp. Fig. 2).

Data Input and Output

We train a separate CNN for each region shown in Supp. Fig. 1 between 70S and 70N. The CNN trained on Outer Asia was used to evaluate pixels within Oceania due to the small sample size of the region. The model used to evaluate Northern Africa was trained on all of Africa but only evaluated over Northern Africa. A separate CNN was trained over Southern Africa. Pixels used

for training are shown in Supp. Fig. 3 (purple shading), and were chosen to optimize computation time while maximizing the representation of unique geophysical features and human impacts.

Over all trained CNNs, we train on 3.8 million HFI pixels and predict 1.4e8 ml-HFI values for 2000 and 2019, separately. We only train on approximately 3% of the non-water data available due to computational and storage limitations, but expect many more samples could be included to further improve the ml-HFI accuracy. Because across the globe there are substantially more low HFI pixels than high pixels (see Fig. 1 and Supp. Fig. 4)), we restrict training to a maximum of 40,000 pixels per 10 deg. by 10 deg. box, and enforce that no HFI bin is represented by more than about 4,000 samples each. Since high HFI values are quite rare (see Supp. Fig. 4), we augment our training set by rotating the inputs by 90 degrees to create 4 unique training/output samples for a single HFI pixel.

Training and Loss Function

The CNN is trained to minimize the mean squared error of the predicted ml-HFI compared to the Williams-HFI. This is done for the year 2000 data (both the Williams-HFI and the Landsat imagery); the CNN is then used to predict the 2019 ml-HFI when no Williams-HFI exists. We have designed the CNN to output a single, continuous value. However, an HFI value of zero holds special meaning, as it implies a pristine landscape with no human influence. To force the CNN to predict the value of exactly zero (rather than just a small number), we re-label HFI values of zero as -0.2 during training when computing the loss function. Because the sigmoid activation function is applied prior to the output of the CNN, the CNN is unable to ever predict a value below 0; thus, this relabeling encourages the CNN to predict values of exactly zero.

As a preprocessing step, Landsat input channels (1,3,4) are standardized, first dividing by 255 (to convert to values between 0 and 1), then subtracting (0.08,0.24,0.12) and dividing by (0.46,0.58,0.47). In this way, we obtain values that are centered around zero and extend to approximately plus/minus 1.0. Throughout our analysis, we noticed that the general “brightness”

of the Landsat imagery can vary in space, and also between year 2000 and 2019 within the GFCv1.7 data. We add a random integer between (-10,10) to each input channel prior to training. In this way, we eliminate the tendency of the CNN to base its predictions too heavily on the magnitude of the input values. This step may be exceptionally important if one wishes to train and predict from various versions of the GFC Landsat imagery.

We train each CNN using a learning rate of 0.001 and a batch size of 128. After selecting a maximum of 40,000 samples for each 10 deg. by 10 deg. Landsat granule, we withhold 750 samples for validation during training. These samples are balanced to distribute similar counts in each of 10 HFI bins. Training is halted via early stopping when the validation mean squared error increases for 3 epochs in a row. At that point, we restore the weights from the best model from prior epochs to define the final CNN.

Attribution of the CNN's Decisions using Layer-wise Relevance Propagation

To gain insight and confidence in the CNN predictions, we apply a technique called layerwise relevance propagation (LRP; (21, 22)). This method highlights the regions of the input, or the pixels of the 3 Landsat images, that were most relevant to the CNNs output (i.e prediction of the ml-HFI). This technique acts to "open the black box" for neural network tasks and, importantly, provides intuitive confirmation or possible refutation of the CNN's decision-making process. In our case, we use LRP as a quality control step to check that the CNN is paying attention to the correct features that would lead to high or low HFI predicted values.

LRP, applied separately to each prediction of the CNN, produces a heatmap displaying the relevance of each pixel of each of the 3 input Landsat images (i.e. channels). We have computed the LRP relevance heatmaps for selected predictions of the 2019 ml-HFI where human pressure has substantially increased. The average relevance across the 3 input channels are displayed in Fig. 2 and in Supp. Fig. 2 where warmer colors denote pixels that were most relevant to the HFI prediction of the CNN. For all of these examples, the CNN places the highest relevance on features within the image that are clearly due to human activity as seen by eye in the 2019

Landsat imagery. This provides us with confidence that the CNN predictions are correct for the right reasons.

Evaluation/Performance

We assess the performance of the year-2000 ml-HFI relative to the year-2000 Williams-HFI by examining both the accuracy (Supp. Fig. 5) and absolute error (Supp. Fig. 6). The ml-HFI values are within an error tolerance of 0.1 of the Williams-HFI values for 73% of the terrestrial globe. When the error tolerance is increased to 0.2, the percentage jumps to 93%. Mean absolute errors as a function of the Williams-HFI values (binned in increments of 0.1) are shown in Supp. Fig. 6. The median absolute error is well below 0.1 for most bins, and this is especially true for regions of very low or very high HFI. The training samples and testing samples generate similar error distributions, which suggests that the CNN models were properly trained. Global maps of the errors and absolute errors are displayed in Supp. Fig. 7-8. Averaged over the entire globe, the ml-HFI has an absolute error of 0.07 and root-mean-squared-error of 0.102 when compared to the Williams-HFI.

To demonstrate the success of the ml-HFI in predicting human pressure from GFCv1.7 Landsat imagery, Supp. Figs. 9-12 provide example outputs of the ml-HFI for urban, agricultural, mining, and rural landscapes in the year-2000. In all of these cases, the ml-HFI compares well with the Williams-HFI, although it is clear that the Williams-HFI emphasizes road networks and is smoother than the ml-HFI which is predicted on a pixel-by-pixel basis.

We assess the skill in the ml-HFI values using the Williams-HFI values as “truth”; however, the Williams-HFI values themselves are not ground truth, but are estimations that can also be incorrect at times (although this is likely rare due to the extensive amount of work that goes into developing this index). With that said, we have found multiple locations where the ml-HFI predicts high human pressure but the Williams-HFI does not. Upon inspection of these cases believe the ml-HFI to be correct. Three such examples, provided in Supp. Fig. 13, highlight how the ml-HFI may be useful for identifying regions of high human pressure that may have been

overlooked in the Williams-HFI.

There are also regions where we believe the ml-HFI struggles. The largest errors appear over dry, rocky terrain devoid of human activity, for example, the rocky parts of the Saharan Desert, the Australian desert, and the spine of the Andes (Supp. Fig. 7-8). In these areas, the ml-HFI predicts human activity where none exists. We provide three such examples in Supp. Fig. 14. However, ml-HFI errors over these problematic regions of low HFI are still relatively small (Supp. Fig. 6-8), with the 90th percentile error being approximately 0.12.

Calculation and Presentation of HFI Changes between 2000 and 2019

We compute the change in ml-HFI between 2000 and 2019 and plot regions with exceptionally large changes (changes greater than 0.25 are shown in red and changes less than -0.25 are shown in blue) on top of the year-2000 ml-HFI values in Fig 2, which are shaded in gray. Because the mI-HFI appears to struggle over rocky deserts and mountains where the Williams-FHI is zero, we mask those change calculations. This masking is carried out in regions where the Williams-HFI was less than 0.02 in 2000 (termed "Wilderness" by (10)) and the mI-HFI error was larger than 0.001. In other words, we mask out changes where the 2000 mI-HFI initially struggled with a zero prediction (i.e. no human activity at all).

Green Outlines in Figure 2

The green outlines in Fig 2 summarize countries that are experiencing an increase in HFI between 2000 and 2019 of greater than 0.25 across 1% or more of their land area and are 'moderately improving' or 'on track to achieve SDG15' (Supp Table 1). All 43 countries in the table have been extensively quality-controlled and verified to have experienced change using the time-lapse feature on Google Earth Pro. Norway, Gabon, Libya and Sudan have been removed from the list as further inspection shows that the CNN struggles in those locations. In Norway the presence of high latitude rocky areas are mis-classified as experiencing change. In Gabon there is an apparent Landsat imagery issue which appears to be localized specifically to Gabon's interior. Both Libya and Sudan contain large areas of rocky, remote deserts that are misclassified

as experiencing substantial human impact.

Acknowledgments

This research was funded, in part, by the United States National Aeronautics and Space Administration (NASA) under grant #18-SLSCVC18-0006.

Data Availability and Code Availability

The ml-HFI for 2000 and 2019 as well as the weights for the trained convolutional neural networks will be made available to the community via the Mountain Scholar permanent data repository with a permanent DOI and via github. All data used in this study is publically available. Human Footprint Index data were downloaded from the github repository https://github.com/scabecks/humanfootprint_2000-2013. Sustainable Development Goal data were downloaded from the United Nations website https://dashboards.sdgindex.org/downloads. Landsat imagery was downloaded from the Hansen Global Forest Change v1.7 repository https://earthenginepartners.appspot.com/science-2013-global-forest/download_v1.7.html .

References

1. O. Venter, *et al.*, Sixteen years of change in the global terrestrial human footprint and implications for biodiversity conservation. *Nat. Commun.* **7**, 12558 (2016).

2. J. Riggio, *et al.*, Global human influence maps reveal clear opportunities in conserving Earth's remaining intact terrestrial ecosystems. *Glob. Chang. Biol.* (2020) https:/doi.org/10.1111/gcb.15109.

3. M. Di Marco, O. Venter, H. P. Possingham, J. E. M. Watson, Changes in human footprint drive changes in species extinction risk. *Nat. Commun.* **9**, 4621 (2018).

4. J. E. M. Watson, *et al.*, The exceptional value of intact forest ecosystems. *Nat Ecol Evol* **2**, 599–610 (2018).

5. R. T. Belote, *et al.*, Mammal species composition reveals new insights into Earth's

remaining wilderness. *Front. Ecol. Environ.* **7**, 10399 (2020).

6. H. L. Beyer, O. Venter, H. S. Grantham, J. E. M. Watson, Substantial losses in ecoregion intactness highlight urgency of globally coordinated action. *CONSERVATION LETTERS* **13**, 1 (2020).

7. U. Pascual, *et al.*, Valuing nature's contributions to people: the IPBES approach. *Current Opinion in Environmental Sustainability* **26-27**, 7–16 (2017).

8. M. H. Ruckelshaus, *et al.*, The IPBES Global Assessment: Pathways to Action. *Trends Ecol. Evol.* **35**, 407–414 (2020).

9. O. Venter, *et al.*, Global terrestrial Human Footprint maps for 1993 and 2009. *Sci Data* **3**, 160067 (2016).

10. B. A. Williams, *et al.*, Change in terrestrial human footprint drives continued loss of intact ecosystems. *bioRxiv*, 2020.05.04.077818 (2020).

11. J. Sachs, G. Schmidt-Traub, G. Lafortune, Speaking truth to power about the SDGs. *Nature* **584**, 344 (2020).

12. R. Naidoo, B. Fisher, Reset Sustainable Development Goals for a pandemic world. *Nature* **583**, 198–201 (2020).

13. , Time to revise the Sustainable Development Goals. *Nature* **583**, 331–332 (2020).

14. J. Sachs, *et al.*, "The Sustainable Development Report 2020 - The Sustainable Development Goals and COVID-19" (Cambridge University Press, 2020).

15. Government of the Co-operative Republic of Guyana, "Guyana First Voluntary National Review of the Sustainable Development Goals" (Ministry of Finance, 2019).

16. A. Krizhevsky, I. Sutskever, G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks" in *Advances in Neural Information Processing Systems*

25, F. Pereira, C. J. C. Burges, L. Bottou, K. Q. Weinberger, Eds. (Curran Associates, Inc., 2012), pp. 1097–1105.

17. F. M. Hoffman, *et al.*, Data Mining in Earth System Science (DMESS 2011). *Procedia Comput. Sci.* **4**, 1450–1455 (2011).

18. V. Kumar, X. Jia, A. Khandelwal, A. Karpatne, Predicting land covers from satellite images using temporal and spatial contexts. *US Patent* (2019) (April 8, 2020).

19. C. Qiu, M. Schmitt, C. Geiß, T.-H. K. Chen, X. X. Zhu, A framework for large-scale mapping of human settlement extent from Sentinel-2 images via fully convolutional neural networks. *ISPRS J. Photogramm. Remote Sens.* **163**, 152–170 (2020).

20. M. C. Hansen, *et al.*, High-resolution global maps of 21st-century forest cover change. *Science* **342**, 850–853 (2013).

21. S. Bach, *et al.*, On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. *PLoS One* **10**, e0130140 (2015).

22. G. Montavon, S. Lapuschkin, A. Binder, W. Samek, K.-R. Müller, Explaining nonlinear classification decisions with deep Taylor decomposition. *Pattern Recognit.* **65**, 211–222 (2017).

23. High Level Expert Group on AI, "Ethics Guidelines for Trustworthy AI: High-Level Expert Group on Artificial Intelligence" (European Commission, 2019).

24. P. J. Phillips, C. A. Hahn, P. C. Fontana, D. A. Broniatowski, M. A. Przybocki, Four Principles of Explainable Artificial Intelligence (2020) https:/doi.org/10.6028/NIST.IR.8312-draft.

25. R. Biggs, S. R. Carpenter, W. A. Brock, Turning back from the brink: detecting an impending regime shift in time to avert it. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 826–831 (2009).

26. J. C. Rocha, G. Peterson, Ö. Bodin, S. Levin, Cascading regime shifts within and across scales. *Science* **362**, 1379–1383 (2018).

27. S. Sarkar, Wilderness preservation and biodiversity conservation—keeping divergent goals distinct. *Bioscience* **49**, 405–412 (1999).

28. T. Newbold, *et al.*, Has land use pushed terrestrial biodiversity beyond the planetary boundary? A global assessment. *Science* **353**, 288–291 (2016).

29. E. C. Ellis, E. C. Antill, H. Kreft, All is not loss: plant biodiversity in the anthropocene. *PLoS One* **7**, e30535 (2012).

30. P. G. Curtis, C. M. Slay, N. L. Harris, A. Tyukavina, M. C. Hansen, Classifying drivers of global forest loss. *Science* **361**, 1108–1111 (2018).

31. M. Lenzen, *et al.*, International trade drives biodiversity threats in developing nations. *Nature* **486**, 109–112 (2012).

32. M. Lenzen, *et al.*, International trade of scarce water. *Ecol. Econ.* **94**, 78–85 (2013).

33. S. H. M. Butchart, *et al.*, Improvements to the Red List Index. *PLoS One* **2**, e140 (2007).

34. E. C. Ellis, N. Ramankutty, Putting people in the map: anthropogenic biomes of the world. *Front. Ecol. Environ.* **6**, 439–447 (2008).

35. The Republic of Gambia, "The Gambia. Voluntary National Review. A report on the progress of implementation of SDGs" (2020).

36. Royaume du Maroc, "Examen National Volontaire de la Mise en Oeuvre des Objectifs de Développement Durable" (2020).

37. E. W. Sanderson, *et al.*, The Human Footprint and the Last of the Wild: The human footprint is a global map of human influence on the land surface, which suggests that human beings are stewards of nature, whether we like it or not. *Bioscience* **52**, 891–904 (2002).

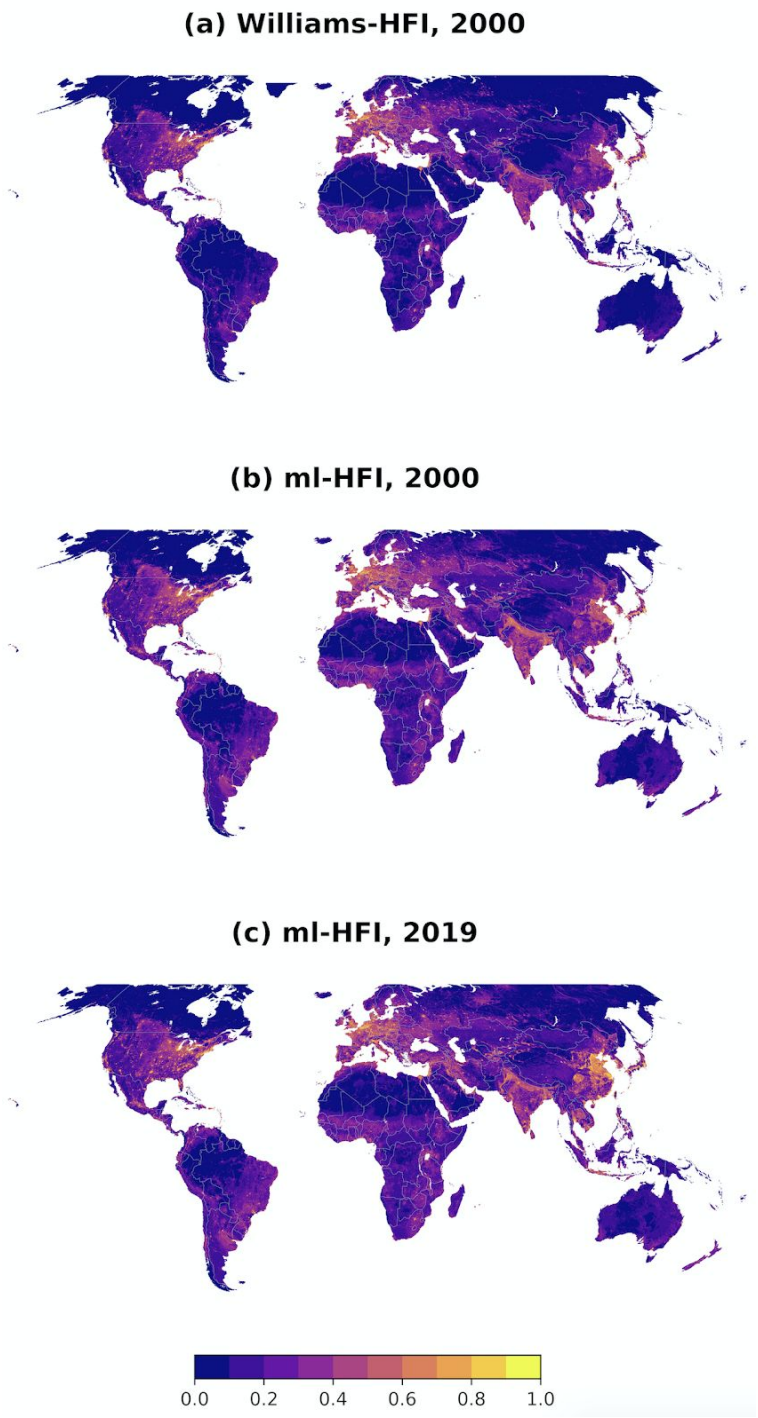

Figure 1: The human footprint index (HFI), a measurement of human pressure on the landscape scaled between 0 and 1. (a) Year-2000 HFI as computed by Williams et al. 2020, (b) year-2000 HFI as computed by a machine learning method, (c) as in (b) but for year-2019.

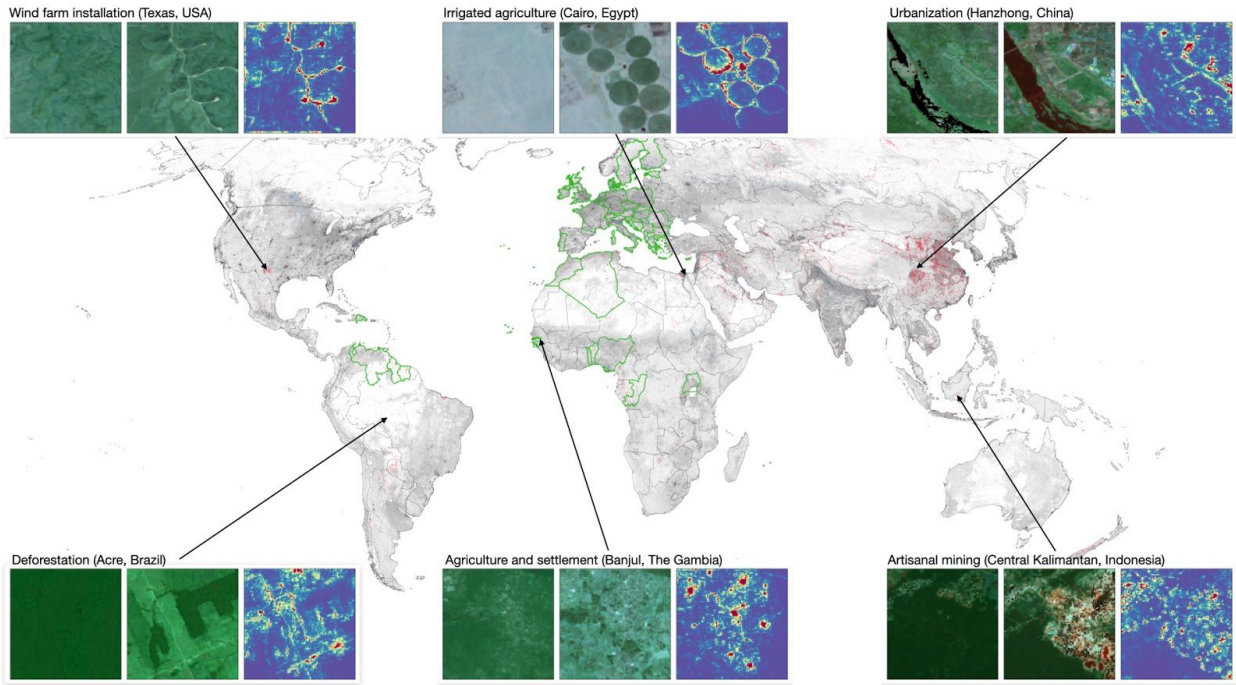

Figure 2: Large changes in human pressure defined as the ml-HFI in 2019 minus that in 2000, where red shading denotes changes larger than 0.25 and blue shading denotes changes smaller than -0.25. Gray shading denotes the ml-HFI from 2000 for reference. Green outlined countries are experiencing substantial increases in HFI and making progress toward SDG15. Inset panels provide examples of increasing human pressure and the relevant features used by the CNN to identify human activity. From left to right, each inset shows (left) year-2000 GFCv1.7 imagery, (middle) year-2019 GFCv1.7 imagery and (right) features most relevant to the CNN for its year-2019 prediction of the ml-HFI. The GFCv1.7 imagery is plotted in false color as its spectral bands are outside of the visible spectrum.

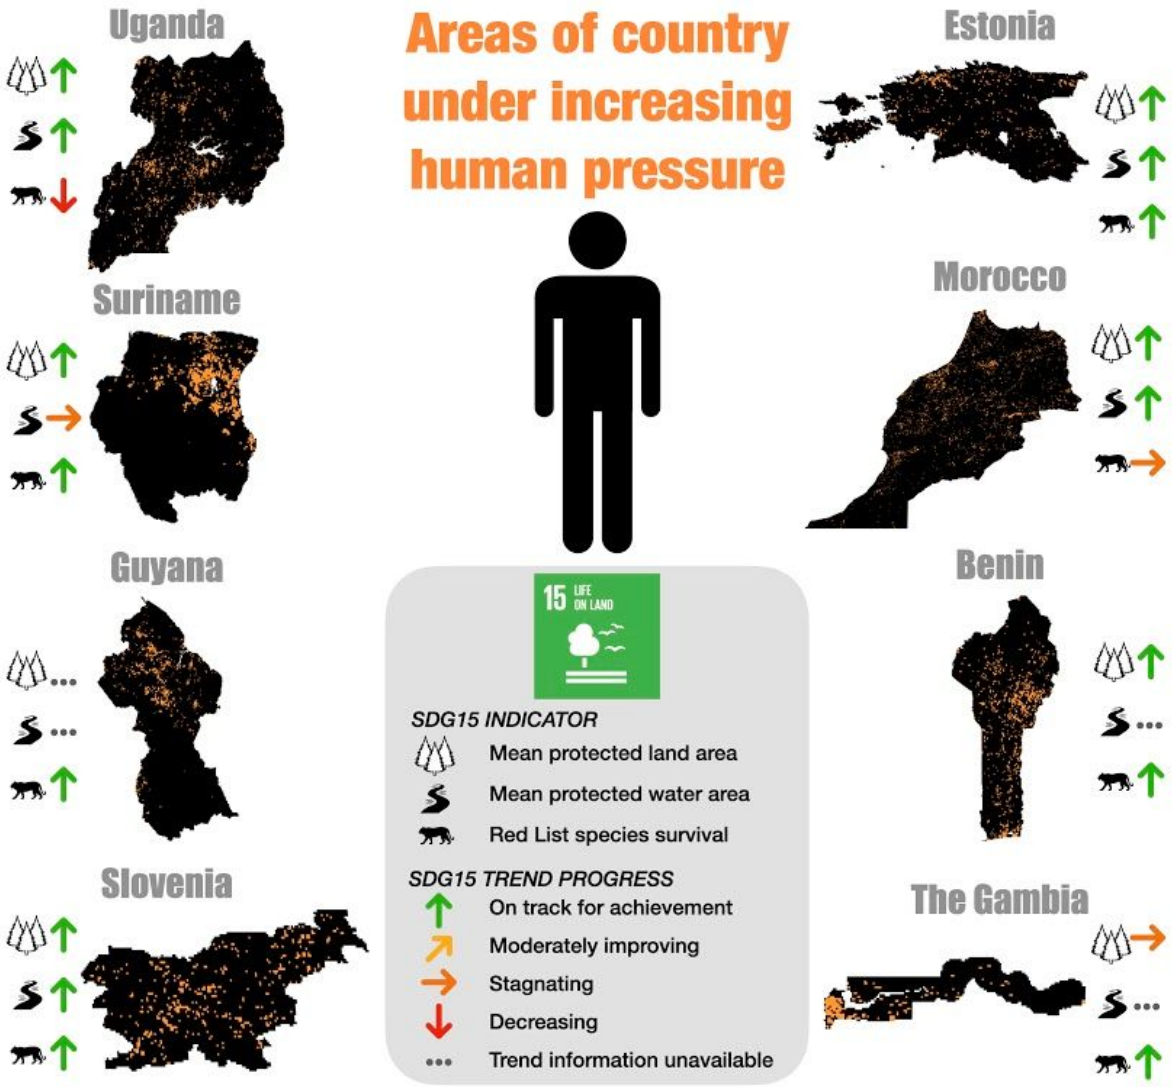

Figure 3: Countries experiencing substantial changes in human pressure that are also making progress toward SDG15. Percent indicates the fraction of each countries' land surface area that is experiencing substantial increases in human pressure, defined as a change of 0.25 on the 0-1 scale of the ml-HFI, over land surface area greater than or equal to 1% of the total country area. Countries shown are a subset of the full list of 43 countries meeting these criteria detailed in Supp Table 1. ml-HFI changes displayed here are degraded in resolution for easier visualization.