Biological Conservation 279 (2023) 109905

Contents lists available at ScienceDirect

Biological Conservation

journal homepage: www.elsevier.com/locate/biocon

Detecting wildlife trafficking in images from online platforms: A test case using deep learning with pangolin images

Ana Sofia Cardoso [a,b,c,*], Sofiya Bryukhova [d], Francesco Renna [e], Luís Reino [a,c], Chi Xu [f], Zixiang Xiao [f], Ricardo Correia [d,g,h], Enrico Di Minin [d,g,i], Joana Ribeiro [a,c,j], Ana Sofia Vaz [a,b,c]

[a] CIBIO, Centro de Investigação em Biodiversidade e Recursos Genéticos, InBIO Laboratório Associado, Campus de Vairão, Universidade do Porto, 4485-661 Vairão, Portugal
[b] Departamento de Biologia, Faculdade de Ciências, Universidade do Porto, 4099-002 Porto, Portugal
[c] BIOPOLIS Program in Genomics, Biodiversity and Land Planning, CIBIO, Campus de Vairão, 4485-661 Vairão, Portugal
[d] Helsinki Lab of Interdisciplinary Conservation Science, Department of Geosciences and Geography, University of Helsinki, Helsinki 00014, Finland
[e] INESC TEC, Faculdade de Ciências da Universidade do Porto, Rua do Campo Alegre s/n, 4169-007 Porto, Portugal
[f] School of Life Sciences, Nanjing University, Nanjing 210023, China
[g] Helsinki Institute of Sustainability Science (HELSUS), University of Helsinki, Helsinki 00014, Finland
[h] Biodiversity Unit, University of Turku, 20014 Turku, Finland
[i] School of Life Sciences, University of KwaZulu-Natal, Durban 4041, South Africa
[j] CIBIO, Centro de Investigação em Biodiversidade e Recursos Genéticos, InBIO Laboratório Associado, Instituto Superior de Agronomia, Universidade de Lisboa, 1349-017 Lisboa, Portugal

ARTICLE INFO

Keywords:
Artificial intelligence
Computer vision
E-commerce
Digital conservation
Pangolin trade
Online trade

ABSTRACT

E-commerce has become a booming market for wildlife trafficking, as online platforms are increasingly more accessible and easier to navigate by sellers, while still lacking adequate supervision. Artificial intelligence models, and specifically deep learning, have been emerging as promising tools for the automated analysis and monitoring of digital online content pertaining to wildlife trade. Here, we used and fine-tuned freely available artificial intelligence models (i.e., convolutional neural networks) to understand the potential of these models to identify instances of wildlife trade. We specifically focused on pangolin species, which are among the most trafficked mammals globally and receiving increasing trade attention since the COVID-19 pandemic. Our convolutional neural networks were trained using online images (available from iNaturalist, Flickr and Google) displaying both traded and non-traded pangolin settings. The trained models showed great performances, being able to identify over 90 % of potential instances of pangolin trade in the considered imagery dataset. These instances included the showcasing of pangolins in popular marketplaces (e.g., wet markets and cages), and the displaying of commonly traded pangolin parts and derivates (e.g., scales) online. Nevertheless, not all instances of pangolin trade could be identified by our models (e.g., in images with dark colours and shaded areas), leaving space for further research developments. The methodological developments and results from this exploratory study represent an advancement in the monitoring of online wildlife trade. Complementing our approach with other forms of online data, such as text, would be a way forward to deliver more robust monitoring tools for online trafficking.

1. Introduction

The unsustainable harvest and trade of wildlife species is among the main drivers of biodiversity change (Scheffers et al., 2019). Thousands of plant and animal species, including live individuals and their derivatives, are over-exploited and illegally traded (Pires and Moreto, 2016). Global wildlife trade generates yearly revenues of ca. 8–12 billion euros, involving hundreds of major cities (mainly from Asia) and millions of individuals (Barber-Meyer, 2010; Rosen and Smith, 2010). Fish and their derived sustenance products, timber, exotic leathers, and furs are examples that can be found in wildlife markets around the globe; many of them being illegally traded and bought for instance to support

* Corresponding author at: CIBIO, Centro de Investigação em Biodiversidade e Recursos Genéticos, InBIO Laboratório Associado, Campus de Vairão, Universidade do Porto, 4485-661 Vairão, Portugal.
E-mail address: sofia.cardoso@cibio.up.pt (A.S. Cardoso).

https://doi.org/10.1016/j.biocon.2023.109905
Received 11 November 2022; Received in revised form 30 December 2022; Accepted 8 January 2023
Available online 20 January 2023

traditional medicine practices and other wildlife commodities (Wyatt, 2021). Parts and derivatives of mammals (e.g., elephant ivory), reptiles (e.g., skin), and birds (e.g., feathers and beaks) are reported as the most commercialized wildlife items globally (Hastie and McCrea-Steele, 2014; Tien Ming et al., 2014).

Wildlife trade has conventionally been done in physical markets (e.g., wet markets); nonetheless, online resources such as social media (e.g., Facebook, Twitter, Instagram, and WeChat) and virtual commercial platforms (e.g., Amazon, eBay, Shopee and Rakuten) have become increasingly popular for both legal and illegal wildlife trade (Davies et al., 2022; Wyatt et al., 2022; Harrison et al., 2016; Feddema et al., 2021b). Facebook, a 2.85 billion users' network, for instance, has become a major global marketplace for illegal wildlife trading, which led this platform to adopt specific measures to forbid the sale of endangered species, such as the creation of alerting "pop up" messages (Davies et al., 2022). Alongside an increased popularity of these global platforms, other e-commerce platforms have also gained visibility in animal trafficking at specific geographic regions (e.g., Taobao in China; Clément et al., 2022). The rise in online trading popularity can be explained by its ease of accessibility as well as limitless, cost-effective, and mostly unsupervised nature (Xu et al., 2020; Kulkarni and Di Minin, 2022). The relative ease with which pictures are showcased online to publicize wildlife trade products also facilitates anonymous interactions between sellers and buyers, thereby increasing the number of potential customers and trade volumes (Yu and Jia, 2015). This makes the tracking of wildlife transactions, and particularly illegal ones, more difficult as it allows for the seller to be more selective and private (Xiao et al., 2017).

With a considerable establishment of wildlife trade in digital platforms, new methods are necessary to survey and monitor potential instances of illegal activity (Di Minin et al., 2019; Siriwat and Nijman, 2020). Artificial intelligence, and more specifically deep learning, has been emerging as a promising field to support the monitoring of events in a wide range of disciplines, including, for instance, the surveillance of diseases and public health issues (e.g., Șerban et al., 2019) or the detection of handguns and security concerns (e.g., Olmos et al., 2018). Deep learning has also become promising for the fast identification of potentially traded species and instances of wildlife trafficking in the web (Di Minin et al., 2019; Xu et al., 2019). When properly calibrated, deep learning models can be highly successful at analysing the content of any digital data from online pages, personal advertisements, e-commerce platforms, and even citizen science or social media networks (Abadi et al., 2016; Lecun et al., 2015). For instance, deep learning models have allowed researchers to analyse wildlife tracking reports using natural language processing of textual contents (Coughlin et al., 2022), to identify and characterize wildlife marketing and sale on social networks via textual content identification (Xu et al., 2020; Feddema et al., 2021a) and to classify images of exotic pet animals advertised for sale on online websites (Kulkarni and Di Minin, 2022). Yet, the potentialities of preexistent and freely available deep learning tools to identify instances of wildlife trade on online images commonly displayed in the web are far from being explored in practice (Di Minin et al., 2019).

In this exploratory study, we aimed to test whether pre-trained and freely available deep learning models can support the identification of potential instances of online wildlife trafficking, using pangolin (Manidae) trade as a test case. Pangolins are among the most trafficked mammal species in the world (Heinrich et al., 2017), with estimates reaching 2.7 million individuals poached every year in central Africa alone (Ingram et al., 2018). Although the commercial trade of the pangolins has been banned under the *Convention on International Trade in Endangered Species of Wild Fauna and Flora* (CITES; Bergesen et al., 2018; Xu et al., 2016), the decline of pangolin populations due to trafficking has not slowed down (Cheng et al., 2017; Nijman et al., 2016; Volpato et al., 2020). In fact, pangolins have received increasing attention particularly since the COVID-19 pandemic (Liu et al., 2020), attracting an increasing number of clients particularly settled in China and Vietnam, as well as in the United States of America and Europe (Heighton and Gaubert, 2021; Omifolaji et al., 2022). Pangolins are traded for meat consumption, as a luxury food, for skin and leather ornamentals, and for curative and spiritual beliefs associated with their scales in traditional medicine, among others (Gimeno-Gilles et al., 2016; Heinrich et al., 2016; Katuwal et al., 2013; Shairp et al., 2016).

Using images shared online, our study aims to investigate the potential of publicly available deep learning models in practice to: (i) support the identification of pangolin species in both natural and anthropogenic (including market) settings; and (ii) identify instances where pangolin parts and derivates are traded online. To do so, we fine-tuned multiple pre-trained deep learning classification and object-detection models using publicly online images of pangolin species. Then, we compared the performance of the models and discuss their added value and limitations for the potential identification of pangolin trafficking. Our study contributes towards the rapidly growing need for automated assessment of online content to support threatened species conservation, as well as wildlife trade identification (Tuia et al., 2022; Isabelle and Westerlund, 2022; Di Minin et al., 2019).

2. Methods

2.1. Methodological framework

To investigate whether pre-trained and freely available deep learning models can support the identification of potential instances of pangolin commerce online, we followed a four-step methodological framework. First, we collected images displaying pangolins, their traded parts and wildlife markets, from different online platforms. After removing duplicated, blurred, or low-resolution images, we manually annotated and labelled each image based on its main content (e.g., showcasing pangolin individuals or parts). From this set of annotated images, we trained and validated five deep learning classification models and three object detection models. Finally, we evaluated the performance of each model to detect pangolins and their traded parts, using different evaluation metrics. All steps were implemented on Google Colaboratory (Colab), a free Jupyter notebook environment from Google Research that allows the execution of Python code entirely in the cloud, requiring no setup to use and providing free access to computing resources, including GPUs (https://colab.research.google.com/notebooks/intro.ipynb). The training, validation and evaluation processes were performed using free and open-source platforms specialized on the training and inference of deep neural networks (Abadi et al., 2016). Specifically, we used Keras (https://keras.io/) and TensorFlow libraries (https://www.tensorflow.org/) for classification models, and TensorFlow API version 1.15 (https://www.tensorflow.org/versions/r1.15/api_docs/python/tf) for object detection models.

2.2. Online imagery data

2.2.1. Data collection

We compiled a dataset of images displaying pangolin species and their commonly traded parts. The eight pangolin species considered in this study were: White-bellied pangolin (*Phataginus tricuspis*), Giant pangolin (*Smutsia gigantea*), Chinese pangolin (*M. pentadactyla*), Philippine pangolin (*M. culionensis*), Sunda pangolin (*M. javanica*), Indian pangolin (*M. crassicaudata),* Temminck's pangolin (*Smutsia temminckii*) and Black-bellied pangolin (*Phataginus tetradactyla*). Additionally, we also included images without pangolins in settings where they are popularly traded (i.e., wet markets in Asia). Images with pangolins and their parts were obtained from the following online databases: iNaturalist (https://www.inaturalist.org/), Flickr (https://www.flickr.com/) and Google images (https://www.google.com/imghp). Image collection from these platforms was performed by searching for general terms on pangolins (e.g., "Pangolin", "Manidae") and the common English name of the eight species of interest (e.g., "Giant pangolin", "Chinese pangolin"; see Table A1 for more details). Searches were performed manually in iNaturalist and Flickr, while for Google images we implemented a Python 3.8.13 (https://docs.python.org/3.8/reference/) script (https://github.com/hardikvasa/google-images-download) to automate the search and scraping of images. Additional images were retrieved by two authors (CX and ZX) through searches on Baidu and Wechat (using the Chinese name for pangolin, "穿山甲") as well as visits to wet markets in China. The content of each image was manually verified to encompass a variety of situations in which pangolins are often presented for trade: live individuals in their natural habitats, dead pangolins in wildlife markets or as a result of seizure incidents, and close-up shots of pangolin derivatives like scales or full pelts.

For the first classification task, we additionally collected images of wet markets without pangolins. This search was also performed on Google images, using general terms associated with wildlife markets (e.g., "wet market"; see Supplementary Material 1). To increase the number of images, a manual search was also performed using the location feature of Google Images, which were set to Asian countries where pangolins are more in demand: Myanmar, Vietnam, China, Thailand, Singapore, Indonesia, and Laos (Cheng et al., 2017; Lim, 2009).

Each image (n = 2634) underwent a manual verification process. Low resolution (i.e., lower than 465 × 257 pixels) and blurred images, not displaying pangolin species nor trade markets were removed, resulting in a total of 2131 images that were retained for labelling and model calibration. For the second classification task, additional images were collected from Google, following the same procedures mentioned before.

2.2.2. *Data labelling*

For the image classification task, each image was manually classified according to its main content, based on a two-level binary classification process. In the first level (L1), the images were classified as "With pangolins" if their content displayed at least one pangolin or their body parts (n = 1075; Fig. 1a). Otherwise, images were classified as "No pangolins" if their content displayed no pangolins in popular sites for their trade (n = 1056; Fig. 1b). For the second level (L2), each image "With pangolins" was further sub-classified as "Pangolins entirely", if their content displayed an individual or multiple pangolins (n = 1014; Fig. 1c), or as "Pangolin parts", if their content displayed specific parts of pangolins (e.g., pangolin scales; n = 307, after adding extra images to counteract the imbalance between classes; Fig. 1d).

For the object detection task, the position and presence of pangolins in each image was manually annotated by one author (ASC; Fig. 1e and f), using the labelImg tool (https://github.com/tzutalin/labelImg).

Fig. 1. Examples of online images included in the two-level classification and object detection task: L1 classification: "With pangolins" (a) versus "No pangolins" (b), L2 classification: "Pangolins entirely" (c) versus "Pangolins parts" (d), object detection: "Pangolins" (e) and (f). The sources of the images displayed in this figure are referred in the supplementary material.

2.3. Model training

2.3.1. Model selection and parametrization

Before the implementation of deep learning models, all images were resized to the same resolution (464 × 256 pixels for the L1 classification and 387 × 163 pixels for the L2 classification) by considering the mean dimensions of the set, and then normalized to the [0,1] range (Na and Fox, 2020).

For the classification task, five freely available convolutional neural networks (CNNs) were selected: VGG16 (Simonyan and Zisserman, 2015), EfficientNetB0, EfficientNetB1 (Tan and Le, 2019), DenseNet121 and DenseNet201 (Huang et al., 2017). These algorithms were selected because of their ease for transfer learning and high performance on similar classification tasks (Curran et al., 2022; Gomez-Donoso et al., 2021). Details on CNN description, parameterization and implementation can be found in the Supplementary material. For model optimization, we used the Adam optimizer algorithm (Kingma and Ba, 2015), and set the batch size to 10 and the number of epochs to 100. The learning rates were chosen from empirical trials over 100 epochs, with 10^{-4} and 10^{-6} showing the best performances. We also implemented an early stop approach, with a patience value of 16 to regularize the model and minimize the loss function (binary cross entropy).

For the object detection task, we used three freely available and pre-trained CNNs: Faster R-CNN ResNet50 (Ren et al., 2015), Faster R-CNN ResNet101 (Ren et al., 2015) and Faster R-CNN Inception v2 (Szegedy et al., 2016). For each network we established 200,000 steps per epoch, a batch size of 1 and a L2 regularization penalty of 10^{-2}. The maximum number of checkpoint files to be evaluated was set to 1. As in every optimization problem, we intend to minimize the loss function, while maximizing the model's performance. The total loss function is usually computed as the sum of the classification (loss of the classifier that determines the type of target object; e.g., binary cross entropy) and localization losses (loss of the regressor that generates a rectangular box to locate the target object; e.g., Mean-Squared Error). We also considered the average inference time per image, which corresponds to the amount of time taken by the models to process a new image and make a prediction. Details on CNN descriptions, parameterization and implementation can be found in the Supplementary Material.

2.3.2. Model training and validation

For both classification tasks (L1 and L2), the performance of the models was evaluated using 5-fold cross validation over the datasets described in Sections 2.2 and 2.3. Namely, each dataset was divided into five subsets, and at each iteration of the 5-fold cross-validation one subset was used to evaluate the models using the performance metrics presented in Section 2.4. The remaining four subsets were used for training (90 % of the images) and validation (10 % of the images).

To enhance the performance of our models and avoid overfitting, we increased the size of the training dataset through a data augmentation approach in the L2 classification model as well as in the object detection models, by applying random transformations to the set of images. Specifically, for the L2 classification models, we used four random transformations per image using the data generator available in Keras (Chollet, 2015): horizontal flip, zoom (range of 0.2), width shift (range of 0.2) and height shift (range of 0.2). For the object detection models, we implemented five random transformations per image using the "data augmentation options" parameter of the TensorFlow configuration file: horizontal flip, image scale, adjust contrast, adjust brightness, and adjust saturation. Both original and transformed images were considered for training, resulting in a final dataset of 3171, for the L2 models and 3428, for the object detection models.

Additionally, to minimize the difference in the number of images among classes (Shahinfar et al., 2020) for the L2 classification task, we implemented a balancing technique during training, in which class weights inversely proportional to the number of samples in each class were applied to the loss function.

2.3.3. Transfer learning

To improve the performance of our models, we applied a transfer learning strategy, which consisted of initializing each CNN with weights of freely available CNNs pre-trained on databases with similar characteristics to our dataset. Specifically, we used CNNs pre-trained with the ImageNet database (http://www.image-net.org/), for the L1 and L2 classification. For the object detection models, we used CNNs pre-trained with the Microsoft Common Objects in Context (MS COCO; https://cocodataset.org/#home) and iNaturalist (https://www.inaturalist.org/) databases (see Supplementary Material for details).

To achieve a good balance between generalization and specificity of image classification for the L1 and L2 classifications, the first CNN model layers were kept frozen during training with transfer learning, while three fully connected layers in VGG16, EfficientNetB0 and EfficientNetB1, and one fully connected layer in DenseNet121 and DenseNet201, were re-trained (fine-tuned) using our training dataset. Before the output layer, we also included an additional dense layer with 128 units and a rectifier linear unit activation function, to enhance the model's adaptation to the classification tasks. Lastly, the output layer was modified to fit a binary classification.

In the object detection models, all the parameters in the configuration files were kept the same as the ones used during the original training of the networks (https://github.com/tensorflow/models/tree/master/research/object_detection/samples/configs), except for the number of classes, which we changed to 1 in order to fit our detection goals (the "Pangolins" class).

2.4. Model performance evaluation

The performance of each classification model (VGG16, EfficientNetB0, EfficientNetB1, DenseNet121 and DenseNet201) was assessed based on commonly adopted classification metrics (Tharwat, 2018; Table 1): accuracy (ACC), specificity (or True Negative Rate: TNR), sensitivity (recall or True Positive Rate: TPR) and F1-score (F1). For the L1 classification, the term positive stands for the presence of pangolins in the images ("With Pangolins"), whereas the negative stands for pangolin absence ("No Pangolins"), while for the L2 classification, the confusion matrices considered "Pangolins parts" images as positives and "Pangolins entirely" as negatives. For both classifications (L1 and L2), the evaluation metrics were calculated as the mean of the performance metrics obtained over the five different folds (see section 2.3.2). Finally, we used a paired samples *t*-test, with a confidence interval of

Table 1
Evaluation metrics considered in the evaluation of classification and object detection models, with respective examples.

Metric	Example
Classification models	
Accuracy	Calculates the level of correctly classified images as "Pangolins entirely" and as "Pangolins parts" by the model.
Specificity	Shows the proportion of correctly classified images as "Pangolins entirely" by the model, in relation to all actual "Pangolins entirely" images.
Sensitivity (or recall)	Shows the proportion of correctly classified images as "Pangolins parts" by the model, in relation to all actual "Pangolins parts" images.
F1-score	Calculates the level of correctly and incorrectly classified images as "Pangolins entirely" and "Pangolins parts" by the model.
Object detection models	
mean Average Precision (mAP)	Indicates the proportion of correctly detected images as "Pangolins" by the model, in relation to all actual and predicted "Pangolins" images.
Average Recall (AR)	Shows the proportion of correctly detected "Pangolins" images by the model, in relation to all actual "Pangolins" images.

0.05 (Hsu and Lachenbruch, 2005) to test for significant differences in classification metrics between each pair of the five CNN models, for both learning rates (10^{-4} and 10^{-6}).

The performance of each object detection model (Faster R-CNN ResNet50, Faster R-CNN ResNet101 and Faster R-CNN Inception v2) was evaluated using the MS COCO detection metrics and the PASCAL VOC detection metrics (Table 1): mean Average Precision (mAP) and Average Recall (AR). The mAP was computed over different Intersection over Union (IoU) thresholds in the case of the models pre-trained on the MS COCO dataset and over a 0.50 IoU threshold in the models pre-trained on the iNaturalist dataset. IoU thresholds represent the ratio between the area of the intersection and the area of the union of the predicted and actual bounding boxes (Rezatofighi et al., 2019). On the other hand, the AR was only calculated for the models pre-trained on the MS COCO dataset, over 1, 10 and 100 detections, as predefined by the transfer learning architecture (see Table A9).

3. Results

3.1. Classifying online images with pangolins versus without pangolins

When discriminating pictures “With pangolins” from those with “No pangolins” (i.e., L1 classification), the considered CNNs showed high performance, with values ranging from 78.11 to 99.53 (Table 2) and few significant differences in their performance metrics ($p < 0.05$; see Table A8 for full results). Overall, the best performance was obtained by VGG16 with a learning rate of 10^{-6}, in terms of accuracy (96.38), specificity (97.55) and F1-score (96.36) values, followed by the EfficientNetB0 with a learning rate of 10^{-4}, which also showed high accuracy (96.20), sensitivity (93.29) and F1-score (96.03). In the case of the VGG16 model, only 4.7 % of the images displaying pangolins were confused by the model as showing no pangolins (false negatives; Fig. 2a). Likewise, from the images showing no pangolins, in 2.4 % of the cases VGG16 incorrectly predicted the class “With pangolins” (false positives; Fig. 2b). A similar pattern was verified for the remaining CNN models (Tables A4 and A5).

3.2. Identifying pangolin parts in online images

When focusing on the classification of images with “Pangolins entirely” versus “Pangolin parts” (i.e., L2 classification; Table 2), the performance of the CNN models significantly differed according to the learning rate and architecture ($p < 0.05$; see Table A8 for full results). EfficientNetB0 with a learning rate of 10^{-4} showed the most satisfactory results of all architectures, in terms of accuracy (97.35), sensitivity (94.88) and F1-score (94.37), followed by DenseNet201 with a learning rate of 10^{-6}, which also showed a high accuracy (95.76), specificity (97.44), and F1-score (90.85). In accordance with what was found for the L1 classification, only a minor proportion of the images displaying “Pangolins entirely” was predicted as “Pangolin parts” by the EfficientNetB0 model (2 %; Fig. 2c). In 5 % of cases in which images showed “Pangolin parts”, the EfficientNetB0 model assumed to be “Pangolins entirely” (Fig. 2d). A similar pattern was observed for the remaining architectures (Tables A6 and A7). Overall, for the L1 classification, the models mostly failed to classify images whose pangolins are in the background, reduced in size or barely visible (e.g., Fig. 2a), while for the L2 classification, the models incorrectly classified images whose pangolins have central and large-sized elements (e.g., in Fig. 2c, the model focused on the tail of the upper pangolin, which has a relatively large size regarding the proportions of the image, causing the rest of the body to be slightly distorted by the perspective).

3.3. Detecting pangolins in online images

The object detection models (Faster R-CNN ResNet101, Faster R-CNN ResNet50 and Faster R-CNN Inception v2) showed different performances depending on the transfer learning weights (MS COCO versus iNaturalist; Table 3). Overall, object detection models pretrained on the MS COCO dataset showed the most satisfactory results, with both Faster R-CNN ResNet101 and Faster R-CNN Inception v2 showing similar performances in terms of mean Average Precision (93.41 and 93.12, respectively).

The object detection model with overall best performance, i.e., Faster R-CNN ResNet101 with MS COCO weights, mostly failed to detect pangolins in different image settings, including images with more than one pangolin, displaying additional pangolin parts, and/or presenting dark colours and emphasized shadows (Fig. 3).

4. Discussion

4.1. Model performance of pangolin identification in online images

In this exploratory study we aimed to investigate, in practice, the potential of publicly available deep learning models to identify potential instances of wildlife trade in the web. Our results show that automated methods based on deep learning models have a high potential for image classification and detection in the fields of ecology and wildlife conservation, thus providing important tools to combat the illicit trade of animals and plants. In our specific test case applied to pangolin species, both image classification and detection models showed great performances, being able to identify over 90 % of potential instances of pangolin trade over the considered imagery set. Of the tested models,

Table 2
Performance metrics for both learning rate scenarios trained for each model (mean ± standard deviation of the five folds). ACC – Accuracy, TPR – Sensitivity, TNR – Specificity and F1 – F1 score. Results are shown for both L1 and L2 classifications. Bold values indicate the best performance results for classification level, learning rate and metric.

Classification models	lr = 10^{-4}				lr = 10^{-6}			
	ACC	TPR	TNR	F_1	ACC	TPR	TNR	F_1
L1 classification: pangolins versus no pangolins								
VGG16	88.13 ± 7.42	78.11 ± 15.57	98.50 ± 1.22	86.28 ± 9.42	**96.38 ± 0.40**	95.25 ± 1.21	**97.55 ± 0.78**	**96.36 ± 0.51**
EfficientNetB0	**96.20 ± 1.93**	**93.29 ± 3.45**	99.06 ± 0.95	**96.04 ± 2.14**	94.04 ± 1.86	93.73 ± 2.06	94.32 ± 3.05	94.05 ± 1.82
EfficientNetB1	95.59 ± 1.12	91.66 ± 2.80	**99.53 ± 0.56**	95.40 ± 1.34	92.26 ± 1.83	90.69 ± 3.39	93.74 ± 1.59	92.13 ± 2.11
DenseNet121	95.96 ± 1.59	92.76 ± 3.63	99.14 ± 0.40	95.80 ± 1.88	94.74 ± 1.13	94.22 ± 2.13	95.30 ± 1.33	94.73 ± 1.28
DenseNet201	95.87 ± 2.48	93.27 ± 5.84	98.37 ± 0.76	95.67 ± 2.87	95.50 ± 1.61	**96. ± 1.57**	94.90 ± 1.91	95.54 ± 1.63
L2 classification: pangolins entirely versus pangolin parts								
VGG16	92.13 ± 1.57	81.21 ± 4.32	95.46 ± 1.99	82.81 ± 2.91	94.02 ± 0.31	85.76 ± 4.79	96.55 ± 1.23	86.92 ± 0.80
EfficientNetB0	**97.35 ± 1.23**	**94.88 ± 2.69**	98.13 ± 1.23	**94.37 ± 2.46**	94.55 ± 1.73	89.33 ± 2.48	96.16 ± 2.14	88.45 ± 3.42
EfficientNetB1	97.12 ± 1.33	92.93 ± 3.71	**98.42 ± 1.27**	93.81 ± 2.64	94.70 ± 1.20	**92.87 ± 2.01**	95.26 ± 1.09	89.11 ± 1.98
DenseNet121	96.37 ± 2.08	92.64 ± 5.35	97.53 ± 2.30	92.35 ± 4.13	95.53 ± 0.77	90.29 ± 3.07	97.13 ± 0.95	90.41 ± 1.30
DenseNet201	97.28 ± 1.45	93.87 ± 2.89	98.32 ± 1.18	94.16 ± 3.02	**95.76 ± 1.17**	90.36 ± 3.54	**97.44 ± 0.80**	**90.85 ± 2.35**

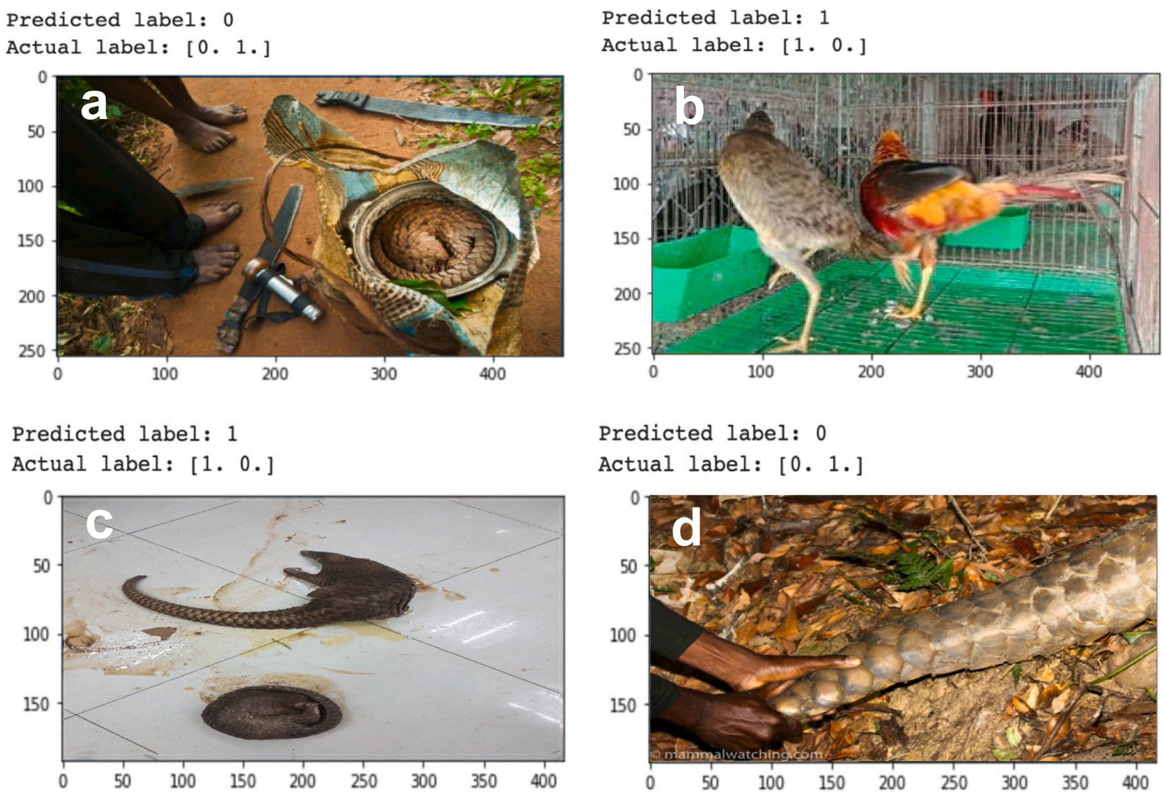

Fig. 2. Examples of images where the classification models failed to predict: (a) pangolins (false negatives) and (b) the absence of pangolins (false positives). The image also shows situations where: (c) "pangolins entirely" were confused by the model as "pangolin parts", or (d) "pangolin parts" were incorrectly classified as "pangolins entirely". The sources of the images displayed in this figure are referred in the supplementary material.

Table 3
Performance metric (mean average precision at an intersection over union threshold of 0.50 [mAP@0.50IOU]), average inference time per image and total loss (sum of the classification and localization losses) for each model. Bold values highlight the best performance and total loss results, as well as the fastest model.

	Faster R-CNN ResNet101 iNaturalist	Faster R-CNN ResNet101 MS COCO	Faster R-CNN ResNet50 iNaturalist	Faster R-CNN Inception v2 MS COCO
Average inference time per image (ms)	395	106	366	**58**
mAP@0.50IOU	92.82	**93.41**	90.10	93.12
Total loss	1.37	**0.69**	1.43	0.81

VGG16, EfficientNetB0, DenseNet201 and Faster R-CNN 101 achieved the most promising results of all architectures. This is in line with previous studies highlighting the performance of VGG16 (Villa et al., 2017; de Silva et al., 2022), EfficientNet (Bodavarapu et al., 2022; Padubidri et al., 2021), DenseNet (Dhillon and Verma, 2022; Jia et al., 2022) and Faster R-CNN ResNet101 (Ibraheam et al., 2021; Peng et al., 2020) in image classification and detection routines for different fields of ecology and wildlife conservation. In our case, whether focusing on pangolins as whole individuals or their sellable parts, the classification models showed promisingly accurate performances. Additionally, such models were also able to accurately identify pangolin individuals in both wild and anthropogenic (including market) settings, as well as instances where pangolin parts and derivates are presented online for trade. These results support the contention that deep learning models can be used to automatically classify online images of endangered species and their derivatives, supporting the identification of wildlife trade situations (Kulkarni and Di Minin, 2021; Xu et al., 2019).

Regarding the performance of individual models, VGG16 with a learning rate of 10^{-6} (L1 classification), EfficientNetB0 with a learning rate of 10^{-4} (L2 classification) and Faster R-CNN ResNet101 pre-trained on the MS COCO dataset (object detection) performed better than the remaining architectures. Deep architectures associated with VGG16 and EfficientNetB0 are particularly robust and easy to implement and retrain, as they replace large kernel-sized filters with smaller ones, reducing the complexity of the network and requiring the training of fewer parameters (Man et al., 2020). EfficientNetB0, specifically, uses compound scaling, a technique that efficiently scales neural networks to accommodate more computational resources. The Faster R-CNN ResNet101 model, in turn, as the name suggests, is a very fast architecture that uses residual connections and batch normalization to extract features at deeper levels. All these characteristics were probably the main drivers of the accurate results observed for VGG16, EfficientNetB0 and Faster R-CNN ResNet101.

Going into more detail, when focusing on the results for the best classification (VGG16 with a learning rate of 10^{-6} and EfficientNetB0 with a learning rate of 10^{-4}) and object detection (Faster R-CNN ResNet101 with MS COCO weights) models, we observed accurate and high performances in classifying and detecting pangolin species and their derivates. However, although both classification and object detection models constitute efficient tools to support the identification of potential situations of pangolin trafficking, classification models achieved slightly higher performances. This may well be attributed to the transfer learning procedure adopted when implementing the classification models, which can improve image classifications of online platforms to produce more robust and reliable identification of pangolins and their derivates in different scenarios. Considering that object detection models required a greater hyperparameter tuning to achieve

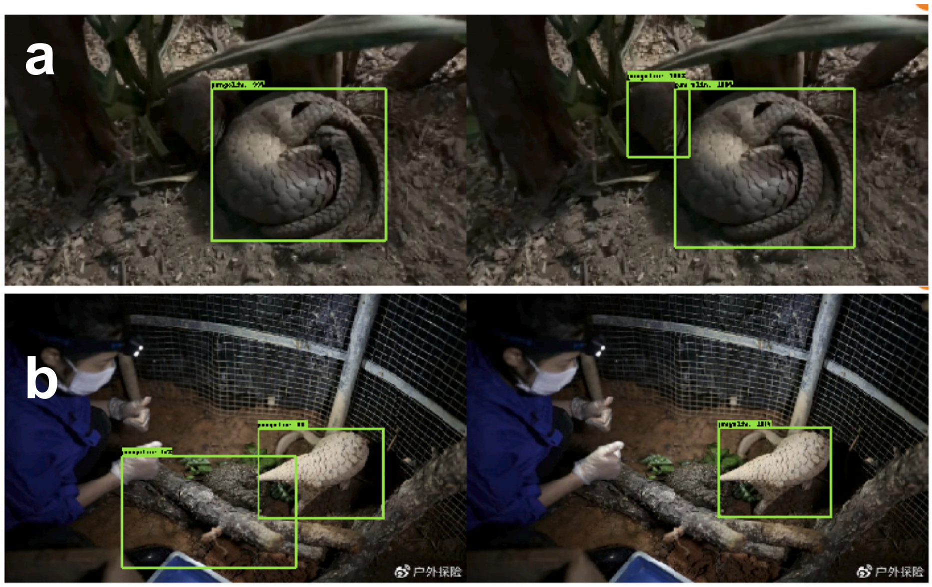

Fig. 3. Examples of images where the Faster R-CNN ResNet101 with MS COCO weights failed to design and predict the pangolins object detection boxes (left images – detected boxes; right images – real boxes). (a) with more than one pangolin and displaying additional pangolin parts; (b) presenting dark colours and emphasized shadows. The sources of the images displayed in this figure are referred in the supplementary material.

stable and satisfactory evaluation results, our results might suggest the preferable use of classification models to track potential instances of pangolin trade efficiently and with less complexity.

4.2. Limitations and the way forward

While the preliminary results of this study are promising, the proposed methodology holds some limitations that need to be tackled before being applied to real-life contexts. Overall, our classification models showed a low misclassification rate when focusing on the identification of instances with and without pangolins, or of pangolin individuals and their sellable parts. However, the models still incorrectly misclassified some images, particularly those where pangolins are in the background, have central and large-sized elements, are reduced in size, are barely visible, or contain dark colours and emphasized shadows. This raises important potential issues in situations where pangolins are present and need to be identified/detected (e.g., when used by online platforms and law enforcement agencies to track and avert online trafficking). In this regard, it is crucial that the models should minimize, as much as possible, the number of false negatives, in order to avoid losing images that contain pangolins. A way forward may be to continue exploring other architectures, such as NasNet (Zoph et al., 2018), GoogLeNet (Szegedy et al., 2015) or Inception-ResNet (Szegedy et al., 2017), as well as additional techniques for model performance improvement (e.g., Cluster-Based Over Sampling). Although we obtained some false negatives, our approach and models still save a huge amount of work and time, especially when considering a large dataset, by avoiding the manual checking of the images. Nevertheless, the results obtained in this study strongly suggest that the combination of manual classification, freely available deep learning architectures, and pre-trained datasets, have great potential to identify pangolins in different contexts.

Additionally, although we collected data closely resembling wildlife trafficking (e.g., pangolins in cages and wet markets), it would be important to collect data from a wider range of online platforms, including e-commerce (e.g., Amazon and eBay) and other social media applications (e.g., Facebook and Twitter). However, the main websites of online wildlife trade (i.e., sites on the dark web) are not accessible by common search engines, which constitutes a major barrier for obtaining more realistic data (Roberts and Hernandez-Castro, 2017). There are also some limitations regarding the application of user-generated content for tracking wildlife trafficking, such as the fact that people may take and post pictures casually, without any kind of illegal intent, making the identification/detection of pangolins by itself insufficient to identify criminal activity. Considering that social media networks typically resize and recompress images with their own preferred settings, a great challenge becomes to reach efficiency and accuracy of computer vision models over social media images with lower resolution and quality. Therefore, incorporating other forms of online data, such as textual information (e.g., tags, captions, comments) may provide a way forward to improve the generalization of deep learning tools for pangolin identification and classification (Feddema et al., 2021a; Kulkarni and Di Minin, 2021; Xu et al., 2019). To note is the fact that social media platforms and other online websites offer a privacy option setting for the public, so that sellers can restrict access to their information and content only to trusted customers (Hastie and McCrea-Steele, 2014).

Considering the high quality of these preliminary results, we hope to contribute towards the development of a more efficient, low cost and less time-consuming tool for supporting the identification and tracking of potential situations of online illegal wildlife sales. Further work is still needed to adapt our models to encompass other taxonomic groups (such as reptiles and birds), which are also known to be widely traded online, though requiring considerable effort at the species-level. Despite these caveats, our exploratory study contributes towards the immense potential of deep learning not only to combat wildlife trade, but also to address wider issues pertaining to instance of wildlife tourism (e.g., in social media platforms) or invasive species (e.g., in citizen science platforms) which are cornerstones subjects in conservation biology.

5. Conclusion

The rise in popularity of online resources and social media platforms has boosted illegal wildlife trade, facilitated by their ease of use, lack of supervision and the possibility to use them anonymously. To support the monitoring and prevention of illegal wildlife sales, we tested in practice the capacity of freely available artificial intelligence models to identify pangolins featuring in online images under multiple settings. Our results showed great performances of the considered classification (e.g., VGG16, EfficientNetB0) and object detection (e.g., Faster R-CNN

ResNet101) models when identifying images with or without pangolins as well as pangolin derivates and sellable parts. It is important to note that further research is needed to robustly identify pangolin trafficking and associated illegal activity. Nonetheless, our approach shows significant potential to help identify online images of pangolins, supporting the recognition of online media advertising wildlife trading. Our exploratory study contributes to an increasingly expanding field of digital conservation where the combination of online data and well-calibrated deep learning models can serve as an important tool to identify content pertaining to wildlife trade, and hence provide more evidence as to the extent of certain species, whether they are protected or not, are involved in online trade. Despite the roads ahead, our proposed methodology is expected to be of interest to other conservation scientists advancing methods for automated identification of human-nature interactions from images (following e.g., Kulkarni and Di Minin, 2022).

CRediT authorship contribution statement

Ana Sofia Cardoso: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Sofiya Bryukhova:** Conceptualization, Software, Investigation, Resources, Data curation. **Francesco Renna:** Software, Writing – original draft, Visualization, Supervision. **Luís Reino:** Writing – original draft, Supervision. **Chi Xu:** Resources, Writing – original draft. **Zixiang Xiao:** Resources, Writing – original draft. **Ricardo Correia:** Writing – original draft, Visualization. **Enrico Di Minin:** Writing – original draft, Visualization. **Joana Ribeiro:** Writing – original draft, Visualization. **Ana Sofia Vaz:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

ASC is supported by the FCT - Portuguese Foundation for Science and Technology through the 2021 PhD Research Studentships [grant reference 2021.05426.BD]. SB and EDM thank the European Research Council (ERC) for funding under the European Union's Horizon 2020 research and innovation program (grant agreement 802933). JR is supported by National Funds through FCT - Fundação para a Ciência e a Tecnologia, I.P., in the scope of the project PTDC/BIA-ECO/0207/2020. RC acknowledges personal financial support from the KONE Foundation (Grant #202101976) and the Academy of Finland (Grant #348352). LR is supported by Portuguese National Funds through FCT, public institute (IP), under the Stimulus of Scientific Employment: Individual Support contracts [contract reference CEECIND/00445/2017]. ASV acknowledges support from FCT - Portuguese Foundation for Science and Technology through the program Stimulus for Scientific Employment - Individual Support [contract reference 2020.01175.CEECIND/CP1601/CT0009].

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.biocon.2023.109905.

References

Abadi, M., Barham, P., Jianmin, C., Zhifend, C., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Zheng, X., 2016. TensorFlow: a system for large-scale machine learning. In: Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI '16). https://doi.org/10.1016/0076-6879(83)01039-3.

Barber-Meyer, S.M., 2010. Dealing with the clandestine nature of wildlife-trade market surveys. Conserv. Biol. 24 (4), 918–923. https://doi.org/10.1111/j.1523-1739.2010.01500.x.

Bergesen, H.O., Parmann, G., Thommessen, Ø.B., 2018. Convention on international trade in endangered species of wild fauna and flora (CITES). In: Yearbook of International Cooperation on Environment and Development 1998–99. Routledge, pp. 156–157.

Bodavarapu, P.N.R., Narayan, T.A., Srinivas, P.V.V.S., 2022. An optimized convolutional neural network model for wild animals detection using filtering techniques and different opacity levels. In: Intelligent Data Communication Technologies and Internet of Things. Springer, Singapore, pp. 1–15.

Cheng, W., Xing, S., Bonebrake, T.C., 2017. Recent pangolin seizures in China reveal priority areas for intervention. Conserv. Lett. 10 (6), 757–764. https://doi.org/10.1111/conl.12339.

Chollet, F., 2015. Keras. GitHub. https://github.com/fchollet/keras.

Clément, T., Haysom, S., Pay, J., 2022. Online markets for pangolin-derived products. In: Dynamics of E-commerce Platforms. Global Initiative Against Transnational Organized Crime, Geneva, Switzerland.

Coughlin, D., Gagnon, M., Grasso, V., Mou, G., Lee, K., Konrad, R., Gore, M., 2022. Extracting and Visualizing Wildlife Trafficking Events From Wildlife Trafficking Reports. arXiv preprint arXiv:2207.08217.

Curran, B., Nekooei, S.M., Chen, G., 2022. Accurate New Zealand wildlife image classification-deep learning approach. In: Australasian Joint Conference on Artificial Intelligence. Springer, Cham, pp. 632–644. https://doi.org/10.1007/978-3-030-97546-3_51.

Davies, A., D'Cruze, N., Senni, C., Martin, R.O., 2022. Global Ecology and Conservation, 33.

Dhillon, A., Verma, G.K., 2022. Fusion of deep learning features for wild animal detection. In: Advances in Deep Learning, Artificial Intelligence and Robotics. Springer, Cham, pp. 15–24.

Di Minin, E., Fink, C., Hiippala, T., Tenkanen, H., 2019. A framework for investigating illegal wildlife trade on social media with machine learning. Conserv. Biol. 33 (1), 210.

Feddema, K., Harrigan, P., Wang, S., 2021a. The dark side of social media engagement: an analysis of user-generated content in online wildlife trade communities. Australas. J. Inf. Syst. 25 https://doi.org/10.3127/ajis.v25i0.2987.

Feddema, K., Nekaris, K.A.I., Nijman, V., Harrigan, P., 2021b. Re-evaluating the notion of value in wildlife trade research from a service marketing perspective. Biol. Conserv. 256, 109060 https://doi.org/10.1016/j.biocon.2021.109060.

Gimeno-Gilles, C., Lelièvre, E., Viau, L., Malik-Ghulam, M., Ricoult, C., Niebel, A., Leduc, N., Limami, A.M., Schmidt-Lebuhn, A.N., Fuchs, J., Hertel, D., Hirsch, H., Toivonen, J., Kessler, M., Loureiro, J., Lavania, U.C.U.C., Srivastava, S., Lavania, S., Basu, S., Yao, J.L., 2016. The conservation status, illegal trade and use of pangolins (Manis spp.). Euphytica 18 (2), 22280.

Gomez-Donoso, F., Escalona, F., Pérez-Esteve, F., Cazorla, M., 2021. Accurate multilevel classification for wildlife images. Comput. Intell. Neurosci. 2021 https://doi.org/10.1155/2021/6690590.

Harrison, J.R., Roberts, D.L., Hernandez-Castro, J., 2016. Assessing the extent and nature of wildlife trade on the dark web. Conserv. Biol. 30 (4), 900–904. https://doi.org/10.1111/cobi.12707.

Hastie, J., McCrea-Steele, T., 2014. Wanted - dead or alive: exposing online wildlife trade. International fund for. Anim. Welf. 33. https://s3.amazonaws.com/ifaw-pantheon/sites/default/files/legacy/IFAW-Wanted-Dead-or-Alive-Exposing-Online-Wildlife-Trade-2014.pdf.

Heighton, S.P., Gaubert, P., 2021. A timely systematic review on pangolin research, commercialization, and popularization to identify knowledge gaps and produce conservation guidelines. Biol. Conserv. 256, 109042 https://doi.org/10.1016/j.biocon.2021.109042.

Heinrich, S., Wittmann, T.A., Prowse, T.A.A., Ross, J.V., Delean, S., Shepherd, C.R., Cassey, P., 2016. Where did all the pangolins go? International CITES trade in pangolin species. Glob. Ecol. Conserv. 8, 241–253. https://doi.org/10.1016/j.gecco.2016.09.007.

Heinrich, S., Wittman, T.A., Ross, J.V., Shepherd, C.R., Challender, D.W.S., Cassey, P., 2017. The global trafficking of pangolin: a comprehensive summary of seizures and trafficking routes from 2010-2015. In: Traffic (Issue December). http://static1.1.sqspcdn.com/static/f/157301/27779360/1513337200303/global-pangolin-assessment.pdf.

Hsu, H., Lachenbruch, P.A., 2005. Paired t test. In: Encyclopedia of Biostatistics, 6. https://doi.org/10.1002/0470011815.b2a15112.

Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks Gao. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4700–4708.

Ibraheam, M., Li, K.F., Gebali, F., Sielecki, L.E., 2021. A performance comparison and enhancement of animal species detection in images with various R-CNN models. AI 2 (4), 552–577. https://doi.org/10.3390/ai2040034.

Ingram, D.J., Coad, L., Abernethy, K.A., Maisels, F., Stokes, E.J., Bobo, K.S., Scharlemann, J.P., 2018. Assessing Africa-wide pangolin exploitation by scaling local data. Conserv. Lett. 11 (2), e12389 https://doi.org/10.1111/conl.12389.

Isabelle, D.A., Westerlund, M., 2022. A review and categorization of artificial intelligence-based opportunities in wildlife, ocean and land conservation. Sustainability 14 (4), 1979.
Jia, L., Tian, Y., Zhang, J., 2022. Identifying animals in camera trap images via neural architecture search. Comput. Intell. Neurosci. 2022.
Katuwal, H.B., Neupane, K.R., Adhikari, D., Thapa, S., 2013. Pangolins Trade, Ethnic Importance and Its Conservation in Eastern Nepal. https://doi.org/10.13140/RG.2.2.27446.91201 (Issue May).
Kingma, D.P., Ba, J.L., 2015. Adam: a method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings, pp. 1–15.
Kulkarni, R., Di Minin, E., 2021. Automated retrieval of information on threatened species from online sources using machine learning. Methods Ecol. Evol. 12 (7), 1226–1239. https://doi.org/10.1111/2041-210X.13608.
Kulkarni, R., Di Minin, E., 2022. Towards Automatic Detection of Wildlife Trade Using Machine Vision Models. arXiv preprint arXiv:2205.11324.
Lecun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. Nature 521 (7553), 436–444. https://doi.org/10.1038/nature14539.
Lim, N., 2009. Ecological research and conservation of Sunda Pangolin Manis Javanica in Singapore. In: Proceedings of the Workshop on Trade and Conservation of Pangolins Native to South and Southeast Asia, 30 June - 2 July 2008, Singapore Zoo (Issue July). http://www.traffic.org/species-reports/traffic_species_mammals51.pdf.
Liu, P., Jiang, J.Z., Wan, X.F., Hua, Y., Li, L., Zhou, J., Chen, J., 2020. Are pangolins the intermediate host of the 2019 novel coronavirus (SARS-CoV-2)? PLoS Pathog. 16 (5), e1008421.
Man, R., Yang, P., Xu, B., 2020. Classification of breast cancer histopathological images using discriminative patches screened by generative adversarial networks. IEEE Access 8, 155362–155377. https://doi.org/10.1109/ACCESS.2020.3019327.
Na, B., Fox, G.C., 2020. Object classifications by image super-resolution preprocessing for convolutional neural networks. Adv. Sci. Technol. Eng. Syst. J. 5 (2), 476–483. https://doi.org/10.25046/aj050261.
Nijman, V., Zhang, M.X., Shepherd, C.R., 2016. Pangolin trade in the mong La wildlife market and the role of Myanmar in the smuggling of pangolins into China. Glob. Ecol. Conserv. 5, 118–126. https://doi.org/10.1016/j.gecco.2015.12.003.
Olmos, R., Tabik, S., Herrera, F., 2018. Automatic handgun detection alarm in videos using deep learning. Neurocomputing 275, 66–72. https://doi.org/10.1016/j.neucom.2017.05.012.
Omifolaji, J.K., Hughes, A.C., Ibrahim, A.S., Zhou, J., Zhang, S., Ikyaagba, E.T., Luan, X., 2022. Dissecting the illegal pangolin trade in China: an insight from seizures data reports. Nat. Conserv. 46, 17. https://doi.org/10.3897/natureconservation.45.57962.
Padubidri, C., Kamilaris, A., Karatsiolis, S., Kamminga, J., 2021. Counting sea lions and elephants from aerial photography using deep learning with density maps. Anim. Biotelemetry 9 (1), 1–10.
Peng, J., Wang, D., Liao, X., Shao, Q., Sun, Z., Yue, H., Ye, H., 2020. Wild animal survey using UAS imagery and deep learning: modified faster R-CNN for kiang detection in Tibetan Plateau. ISPRS J. Photogramm. Remote Sens. 169, 364–376. https://doi.org/10.1016/j.isprsjprs.2020.08.026.
Pires, S.F., Moreto, W., 2016. The illegal wildlife trade. In: Choice Reviews Online, 50. https://doi.org/10.5860/choice.50-5912. Issue 11.
Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. Adv. Neural Inf. Proces. Syst. 28.
Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S., 2019. Generalized intersection over union: a metric and a loss for bounding box regression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 658–666.
Roberts, D.L., Hernandez-Castro, J., 2017. Bycatch and illegal wildlife trade on the dark web. Oryx 51 (3), 393–394. https://doi.org/10.1017/S0030605317000679.
Rosen, G.E., Smith, K.F., 2010. Summarizing the evidence on the international trade in illegal wildlife. EcoHealth 7 (1), 24–32. https://doi.org/10.1007/s10393-010-0317-y.
Scheffers, B.R., Oliveira, B.F., Lamb, I., Edwards, D.P., 2019. Global wildlife trade across the tree of life. Science 366 (6461), 71–76. https://doi.org/10.1126/science.aav5327.
Şerban, O., Thapen, N., Maginnis, B., Hankin, C., Foot, V., 2019. Real-time processing of social media with SENTINEL: a syndromic surveillance system incorporating deep learning for health classification. Inf. Process. Manag. 56 (3), 1166–1184. https://doi.org/10.1016/j.ipm.2018.04.011.
Shahinfar, S., Meek, P., Falzon, G., 2020. "How many images do I need?" Understanding how sample size per class affects deep learning model performance metrics for balanced designs in autonomous wildlife monitoring. Ecol. Informatics 57, 101085.
Shairp, R., Veríssimo, D., Fraser, I., Challender, D., Macmillan, D., 2016. Understanding urban demand for wild meat in Vietnam: implications for conservation actions. PLoS ONE 11 (1), 1–14. https://doi.org/10.1371/journal.pone.0134787.
de Silva, E.M., Kumarasinghe, P., Indrajith, K.K., Pushpakumara, T.V., Vimukthi, R.D., de Zoysa, K., de Silva, S., 2022. Feasibility of using convolutional neural networks for individual-identification of wild Asian elephants. Mamm. Biol. 1–11.
Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. In: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings, pp. 1–14.
Siriwat, P., Nijman, V., 2020. Wildlife trade shifts from brick-and-mortar markets to virtual marketplaces: a case study of birds of prey trade in Thailand. J. Asia-Pac. Biodivers. 13 (3), 454–461.
Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 1–9. https://doi.org/10.1109/CVPR.2015.7298594, 07-12-June.
Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2818–2826.
Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A., 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In: Thirty-first AAAI Conference on Artificial Intelligence.
Tan, M., Le, Q.V., 2019. EfficientNet: Rethinking model scaling for convolutional neural networks. In: 36th International Conference on Machine Learning, ICML 2019, 2019-June, pp. 10691–10700.
Tharwat, A., 2018. Classification assessment methods. Appl. Comput. Informatics 17 (1), 168–192. https://doi.org/10.1016/j.aci.2018.08.003.
Tien Ming, L., Sigouin, A., Pinedo-Vasquez, M., Nasi, R., 2014. The Harvest of Wildlife for Bushmeat and Traditional Medicine in East, South and Southeast Asia.
Tuia, D., Kellenberger, B., Beery, S., Costelloe, B.R., Zuffi, S., Risse, B., Berger-Wolf, T., 2022. Perspectives in machine learning for wildlife conservation. Nat. Commun. 13 (1), 1–15. https://doi.org/10.1038/s41467-022-27980-y.
Villa, A.G., Salazar, A., Vargas, F., 2017. Towards automatic wild animal monitoring: identification of animal species in camera-trap images using very deep convolutional neural networks. Ecol. Informatics 41, 24–32.
Volpato, G., Fontefrancesco, M.F., Gruppuso, P., Zocchi, D.M., Pieroni, A., 2020. Baby pangolins on my plate: possible lessons to learn from the COVID-19 pandemic. J. Ethnobiol. Ethnomed. 16 (1), 1–12. https://doi.org/10.1186/s13002-020-00366-4.
Wyatt, T., 2021. Wildlife trafficking: a deconstruction of the crime, victims and offenders. Springer Nature.
Wyatt, T., Miralles, O., Massé, F., Lima, R., da Costa, T.V., Giovanini, D., 2022. Wildlife trafficking via social media in Brazil. Biol. Conserv. 265, 109420.
Xiao, Y., Guan, J., Xu, L., 2017. Wildlife Cyber Crime in China: E-commerce and Social Media Monitoring in 2016, 12.
Xu, L., Guan, J., Lau, W., Xiao, Y., 2016. An overview of pangolin trade in China. In: Traffic Briefing Paper, 10.
Xu, Q., Li, J., Cai, M., Mackey, T.K., 2019. Use of machine learning to detect wildlife product promotion and sales on twitter. Front. Big Data 2, 28. https://doi.org/10.3389/fdata.2019.00028.
Xu, Q., Cai, M., MacKey, T.K., 2020. The illegal wildlife digital market: an analysis of Chinese wildlife marketing and sale on Facebook. Environ. Conserv. 47 (3), 206–212. https://doi.org/10.1017/S0376892920000235.
Yu, X., Jia, W., 2015. Moving targets: tracking online sales of illegal wildlife products in China. In: TRAFFIC Briefing, 1009649 (February 2015), pp. 1–10.
Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V., 2018. Learning transferable architectures for scalable image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8697–8710.