



**HAL**  
open science

# Trade-off between deep learning for species identification and inference about predator-prey co-occurrence: Reproducible R workflow integrating models in computer vision and ecological statistics

Olivier Gimenez, Maëlis Kervellec, Jean-Baptiste Fanjul, Anna Chaine, Lucile Marescot, Yoann Bollet, Christophe Duchamp

## ► To cite this version:

Olivier Gimenez, Maëlis Kervellec, Jean-Baptiste Fanjul, Anna Chaine, Lucile Marescot, et al.. Trade-off between deep learning for species identification and inference about predator-prey co-occurrence: Reproducible R workflow integrating models in computer vision and ecological statistics. *Computo*, 2022, 10.57750/yfm2-5f45 . hal-03647789

**HAL Id: hal-03647789**

**<https://hal.science/hal-03647789>**

Submitted on 20 Apr 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Trade-off between deep learning for species identification and inference about predator-prey co-occurrence: Reproducible R workflow integrating models in computer vision and ecological statistics

O. Gimenez, M. Kervellec, J.-B. Fanjul, A. Chainé, L. Marescot, Y. Bollet, C. Duchamp

2021-08-26

## Abstract

Deep learning is used in computer vision problems with important applications in several scientific fields. In ecology for example, there is a growing interest in deep learning for automatizing repetitive analyses on large amounts of images, such as animal species identification.

However, there are challenging issues toward the wide adoption of deep learning by the community of ecologists. First, there is a programming barrier as most algorithms are written in `Python` while most ecologists are versed in `R`. Second, recent applications of deep learning in ecology have focused on computational aspects and simple tasks without addressing the underlying ecological questions or carrying out the statistical data analysis to answer these questions.

Here, we showcase a reproducible R workflow integrating both deep learning and statistical models using predator-prey relationships as a case study. We illustrate deep learning for the identification of animal species on images collected with camera traps, and quantify spatial co-occurrence using multispecies occupancy models.

Despite average model classification performances, ecological inference was similar whether we analysed the ground truth dataset or the classified dataset. This result calls for further work on the trade-offs between time and resources allocated to train models with deep learning and our ability to properly address key ecological questions with biodiversity monitoring. We hope that our reproducible workflow will be useful to ecologists and applied statisticians.

All material (source of the Rmarkdown notebook and auxiliary files) is available from <https://github.com/oliviergimenez/computo-deeplearning-occupany-lynx>.

## Introduction

Computer vision is a field of artificial intelligence in which a machine is taught how to extract and interpret the content of an image (Krizhevsky, Sutskever, and Hinton 2012). Computer vision relies on deep learning that allows computational models to learn from training data – a set of manually labelled images – and make predictions on new data – a set of unlabelled images (Baraniuk, Donoho, and Gavish 2020; LeCun, Bengio, and Hinton 2015). With the growing availability of massive data, computer vision with deep learning is being increasingly used to perform tasks such as object detection, face recognition, action and activity recognition or human pose estimation in fields as diverse as medicine, robotics, transportation, genomics, sports and agriculture (Voulodimos et al. 2018).

In ecology in particular, there is a growing interest in deep learning for automatizing repetitive analyses on large amounts of images, such as identifying plant and animal species, distinguishing individuals of the same

or different species, counting individuals or detecting relevant features (Christin, Hervet, and Lecomte 2019; Lamba et al. 2019; Weinstein 2018). By saving hours of manual data analyses and tapping into massive amounts of data that keep accumulating with technological advances, deep learning has the potential to become an essential tool for ecologists and applied statisticians.

Despite the promising future of computer vision and deep learning, there are challenging issues toward their wide adoption by the community of ecologists (e.g. Wearn, Freeman, and Jacoby 2019). First, there is a programming barrier as most, if not all, algorithms are written in the `Python` language while most ecologists are versed in `R` (Lai et al. 2019). If ecologists are to use computer vision in routine, there is a need for bridges between these two languages (through, e.g., the `reticulate` package Allaire et al. (2017) or the `shiny` package Tabak et al. (2020)). Second, recent applications of computer vision via deep learning in ecology have focused on computational aspects and simple tasks without addressing the underlying ecological questions (Sutherland et al. 2013), or carrying out statistical data analysis to answer these questions (Gimenez et al. 2014). Although perfectly understandable given the challenges at hand, we argue that a better integration of the *why* (ecological questions), the *what* (automatically labelled images) and the *how* (statistics) would be beneficial to computer vision for ecology (see also Weinstein 2018).

Here, we showcase a full why-what-how workflow in `R` using a case study on the structure of an ecological community (a set of co-occurring species) composed of the Eurasian lynx (*Lynx lynx*) and its two main preys. First, we introduce the case study and motivate the need for deep learning. Second we illustrate deep learning for the identification of animal species in large amounts of images, including model training and validation with a dataset of labelled images, and prediction with a new dataset of unlabelled images. Last, we proceed with the quantification of spatial co-occurrence using statistical models.

## Collecting images with camera traps

Lynx (*Lynx lynx*) went extinct in France at the end of the 19th century due to habitat degradation, human persecution and decrease in prey availability (Vandel and Stahl 2005). The species was reintroduced in Switzerland in the 1970s (Breitenmoser 1998), then re-colonised France through the Jura mountains in the 1980s (Vandel and Stahl 2005). The species is listed as endangered under the 2017 IUCN Red list and is of conservation concern in France due to habitat fragmentation, poaching and collisions with vehicles. The Jura holds the bulk of the French lynx population.

To better understand its distribution, we need to quantify its interactions with its main preys, roe deer (*Capreolus capreolus*) and chamois (*Rupicapra rupicapra*) (Molinari-Jobin et al. 2007), two ungulate species that are also hunted. To assess the relative contribution of predation and hunting, a predator-prey program was set up jointly by the French Office for Biodiversity, the Federations of Hunters from the Jura, Ain and Haute-Savoie counties and the French National Centre for Scientific Research.

Animal detections were made using a set of camera traps in the Jura mountains that were deployed in the Jura and Ain counties (see Figure 1). We divided the two study areas into grids of  $2.7 \times 2.7$  km cells or sites hereafter (Zimmermann et al. 2013) in which we set two camera traps per site (Xenon white flash with passive infrared trigger mechanisms, model Capture, Ambush and Attack; Cuddeback), with 18 sites in the Jura study area, and 11 in the Ain study area that were active over the study period (from February 2016 to October 2017 for the Jura county, and from February 2017 to May 2019 for the Ain county). The location of camera traps was chosen to maximise lynx detection. Camera traps were checked weekly to change memory cards, batteries and to remove fresh snow after heavy snowfall.

In total, 45563 and 18044 pictures were considered in the Jura and Ain sites respectively after manually dropping empty pictures and pictures with unidentified species. Note that classifying empty images could be automatised with deep learning (Norouzzadeh et al. 2021; Tabak et al. 2020). We identified the species present on all images by hand (see Table 1) using `digikam` a free open-source digital photo management application (<https://www.digikam.org/>). This operation took several weeks of labor full time, which is often identified as a limitation of camera trap studies. To expedite this tedious task, computer vision with deep learning has been identified as a promising approach (Norouzzadeh et al. 2021; Tabak et al. 2019; Willi et al. 2019).

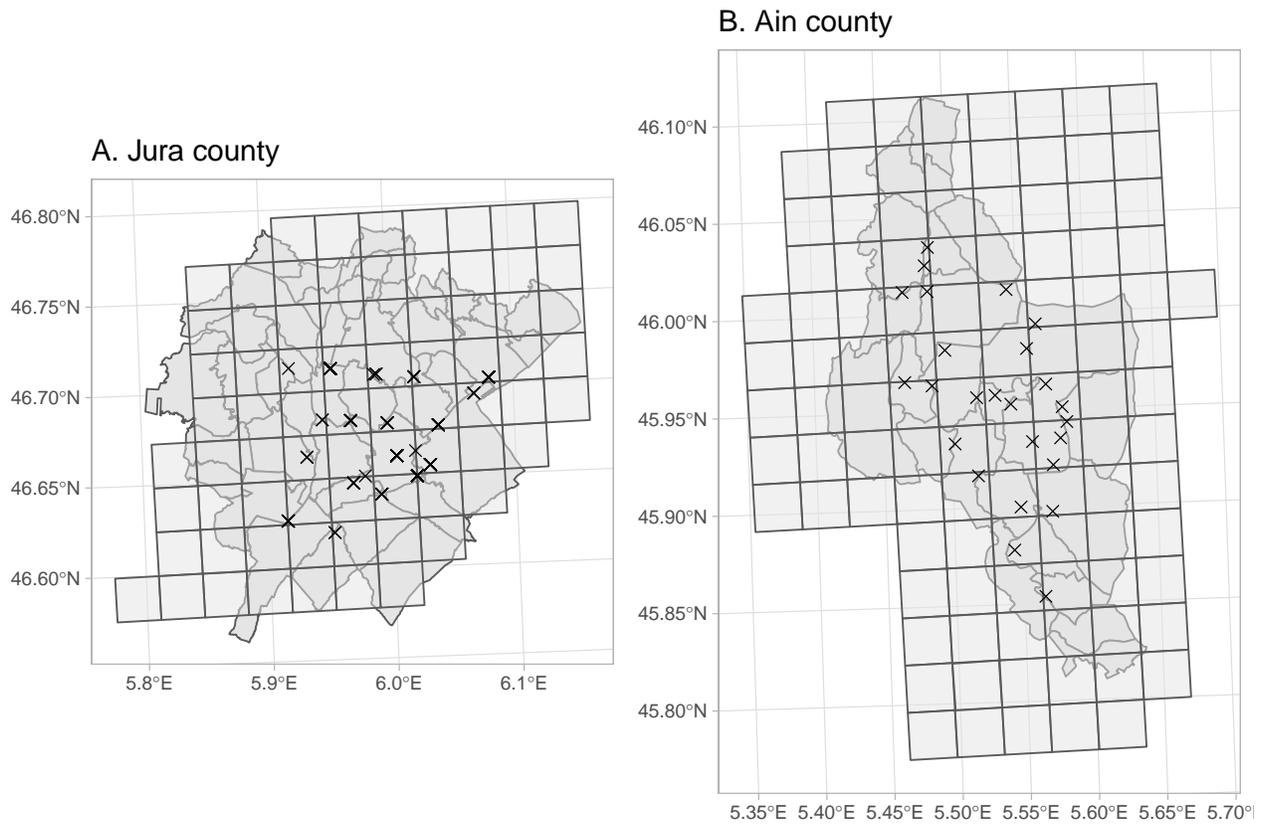


Figure 1: Study area, grid and camera trap locations.

Table 1: Species identified in the Jura and Ain study sites with samples size (n). Only first 10 species with most images are shown.

Species in Jura study site	n	Species in Ain study site	n
human	31644	human	4946
vehicule	5637	vehicule	4454
dog	2779	dog	2310
fox	2088	fox	1587
chamois	919	rider	1025
wild board	522	roe deer	860
badger	401	chamois	780
roe deer	368	hunter	593
cat	343	wild board	514
lynx	302	badger	461

Table 2: Model training performance. Images from the Jura study site were used for training.

species	precision	recall
<b>badger</b>	0.78	0.88
<b>red deer</b>	0.67	0.21
<b>chamois</b>	0.86	0.81
<b>cat</b>	0.89	0.78
<b>roe deer</b>	0.67	0.81
<b>dog</b>	0.78	0.84
<b>human</b>	0.99	0.79
<b>hare</b>	0.32	0.52
<b>lynx</b>	0.87	0.95
<b>fox</b>	0.85	0.90
<b>wild boar</b>	0.93	0.88
<b>vehicule</b>	0.95	0.98

## Deep learning for species identification

Using the images we obtained with camera traps (Table 1), we trained a model for identifying species using the Jura study site as a calibration dataset. We then assessed this model’s ability to automatically identify species on a new dataset, also known as transferability, using the Ain study site as an evaluation dataset.

### Training - Jura study site

We selected at random 80% of the annotated images for each species in the Jura study site for training, and 20% for testing. We applied various transformations (flipping, brightness and contrast modifications; Shorten and Khoshgoftaar (2019)) to improve training (see Appendix). To reduce model training time and overcome the small number of images, we used transfer learning (Yosinski et al. 2014; Shao, Zhu, and Li 2015) and considered a pre-trained model as a starting point. Specifically, we trained a deep convolutional neural network (ResNet-50) architecture (He et al. 2016) using the `fastai` library (<https://docs.fast.ai/>) that implements the `PyTorch` library (Paszke et al. 2019). Interestingly, the `fastai` library comes with an R interface (<https://eagerai.github.io/fastai/>) that uses the `reticulate` package to communicate with `Python`, therefore allowing R users to access up-to-date deep learning tools. We trained models on the Montpellier Bioinformatics Biodiversity platform using a GPU machine (Titan Xp nvidia) with 16Go of RAM. We used 20 epochs which took approximately 10 hours. The computational burden prevented us from providing a full reproducible analysis, but we do so with a subsample of the dataset in the Appendix. All trained models are available from <https://doi.org/10.5281/zenodo.5164796>.

We calculated three metrics to evaluate our model performance at correctly identifying species (e.g. Duggan et al. 2021). Specifically, we relied on *accuracy* the ratio of correct predictions to the total number of predictions, *recall* a measure of false negatives (FN; e.g. an image with a lynx for which our model predicts another species) with  $\text{recall} = \text{TP} / (\text{TP} + \text{FN})$  where TP is for true positives, and *precision* a measure of false positives (FP; e.g. an image with any species but a lynx for which our model predicts a lynx) with  $\text{precision} = \text{TP} / (\text{TP} + \text{FP})$ . In camera trap studies, a strategy (Duggan et al. 2021) consists in optimizing precision if the focus is on rare species (lynx), while recall should be optimized if the focus is on common species (chamois and roe deer).

We achieved 85% accuracy during training. Our model had good performances for the three classes we were interested in, with 87% precision for lynx and 81% recall for both roe deer and chamois (Table 2).

### Transferability - Ain study site

We evaluated transferability for our trained model by predicting species on images from the Ain study site which were not used for training. Precision was 77% for lynx, and while we achieved 86% recall for roe deer,

Table 3: Model transferability performance. Images from the Ain study site were used for assessing transferability.

	precision	recall
<b>badger</b>	0.71	0.89
<b>rider</b>	0.79	0.92
<b>red deer</b>	0.00	0.00
<b>chamois</b>	0.82	0.08
<b>hunter</b>	0.17	0.11
<b>cat</b>	0.46	0.59
<b>roe deer</b>	0.67	0.86
<b>dog</b>	0.77	0.35
<b>human</b>	0.51	0.93
<b>hare</b>	0.37	0.35
<b>lynx</b>	0.77	0.89
<b>marten</b>	0.05	0.04
<b>fox</b>	0.90	0.53
<b>wild board</b>	0.75	0.94
<b>cow</b>	0.01	0.25
<b>vehicule</b>	0.94	0.51

our model performed poorly for chamois with 8% recall (Table 3).

To better understand this pattern, we display the results under the form of a confusion matrix that compares model classifications to manual classifications (Figure 2). There were a lot of false negatives for chamois, meaning that when a chamois was present in an image, it was often classified as another species by our model.

Overall, our model trained on images from the Jura study site did poorly at correctly predicting species on images from the Ain study site. This result does not come as a surprise, as generalizing classification algorithms to new environments is known to be difficult (Beery, Horn, and Perona 2018). While a computer scientist might be disappointed in these results, an ecologist would probably wonder whether ecological inference about the interactions between lynx and its prey is biased by these average performances, a question we address in the next section.

## Spatial co-occurrence

Here, we analysed the data we acquired from the previous section. For the sake of comparison, we considered two datasets, one made of the images manually labelled for both the Jura and Ain study sites pooled together (*ground truth dataset*), and the other in which we pooled the images that were manually labelled for the Jura study site and the images that were automatically labelled for the Ain study site using our trained model (*classified dataset*).

We formatted the data by generating monthly detection histories, that is a sequence of detections ( $Y_{sit} = 1$ ) and non-detections ( $Y_{sit} = 0$ ), for species  $s$  at site  $i$  and sampling occasion  $t$  (see Figure 3).

To quantify spatial co-occurrence between lynx and its preys, we used a multispecies occupancy modeling approach (Rota et al. 2016; Clipp et al. 2021) using the R package `unmarked` (Fiske and Chandler 2011). The multispecies occupancy model assumes that observations  $y_{sit}$ , conditional on  $Z_{si}$  the latent occupancy state of species  $s$  at site  $i$  are drawn from Bernoulli random variables  $Y_{sit}|Z_{si} \sim \text{Bernoulli}(Z_{si}p_{sit})$  where  $p_{sit}$  is the detection probability of species  $s$  at site  $i$  and sampling occasion  $t$ . Detection probabilities can be modeled as a function of site and/or sampling covariates, or the presence/absence of other species, but for the sake of illustration, we will make them only species-specific here.

The latent occupancy states are assumed to be distributed as multivariate Bernoulli random variables

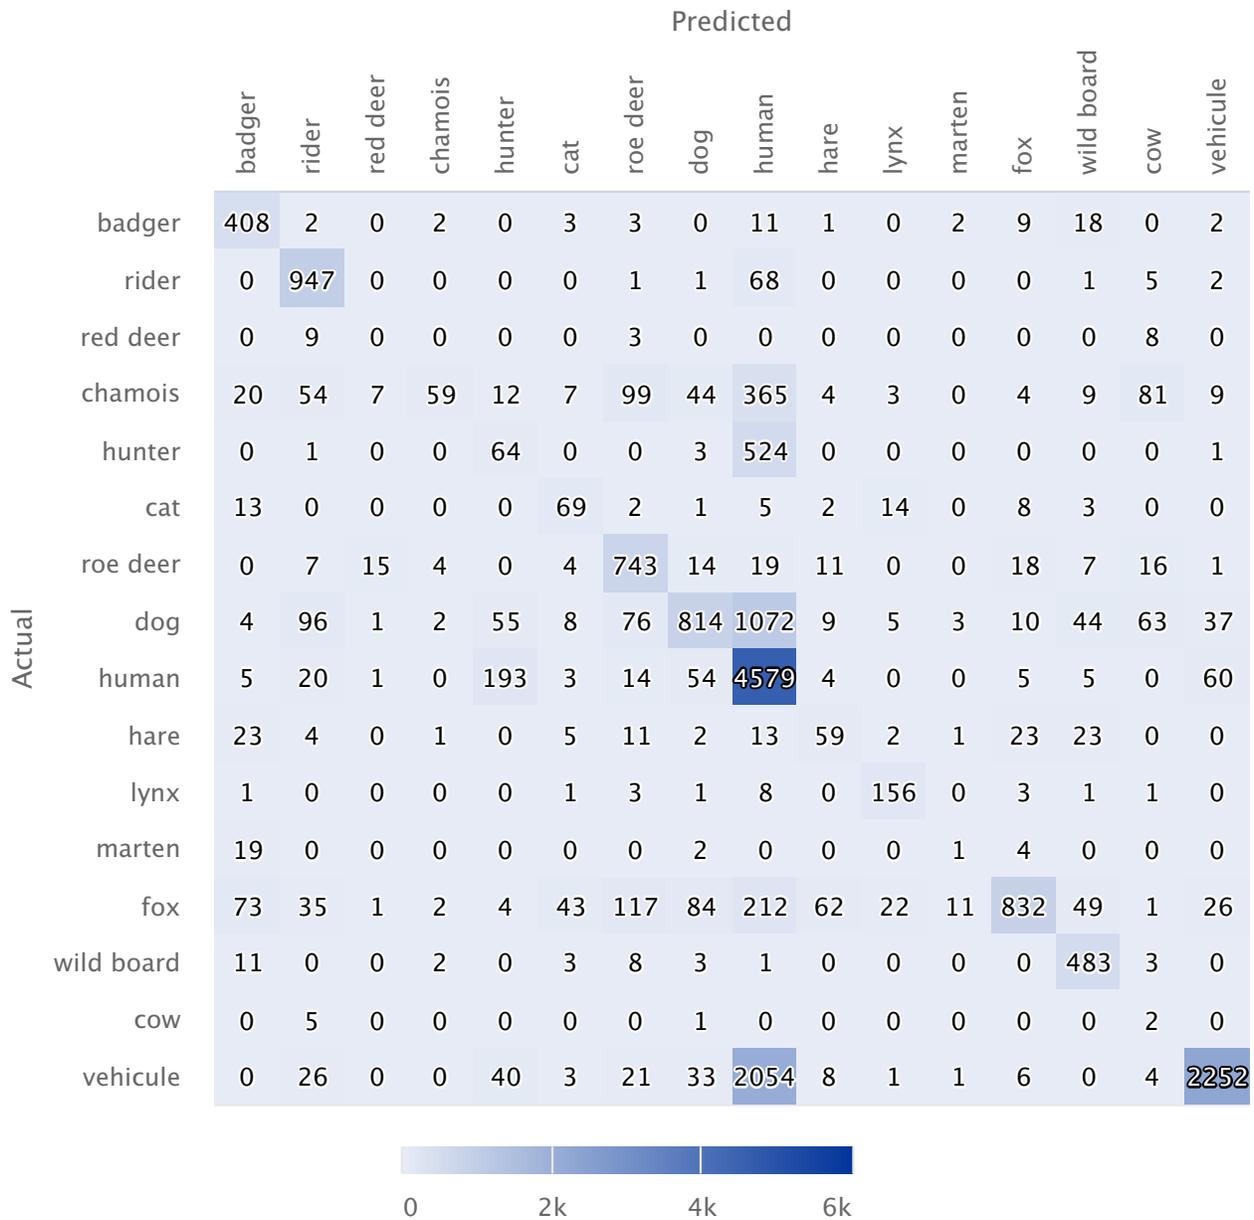


Figure 2: Confusion matrix comparing automatic to manual species classifications. Species that were predicted by our model are in columns, and species that are actually in the images are in rows.

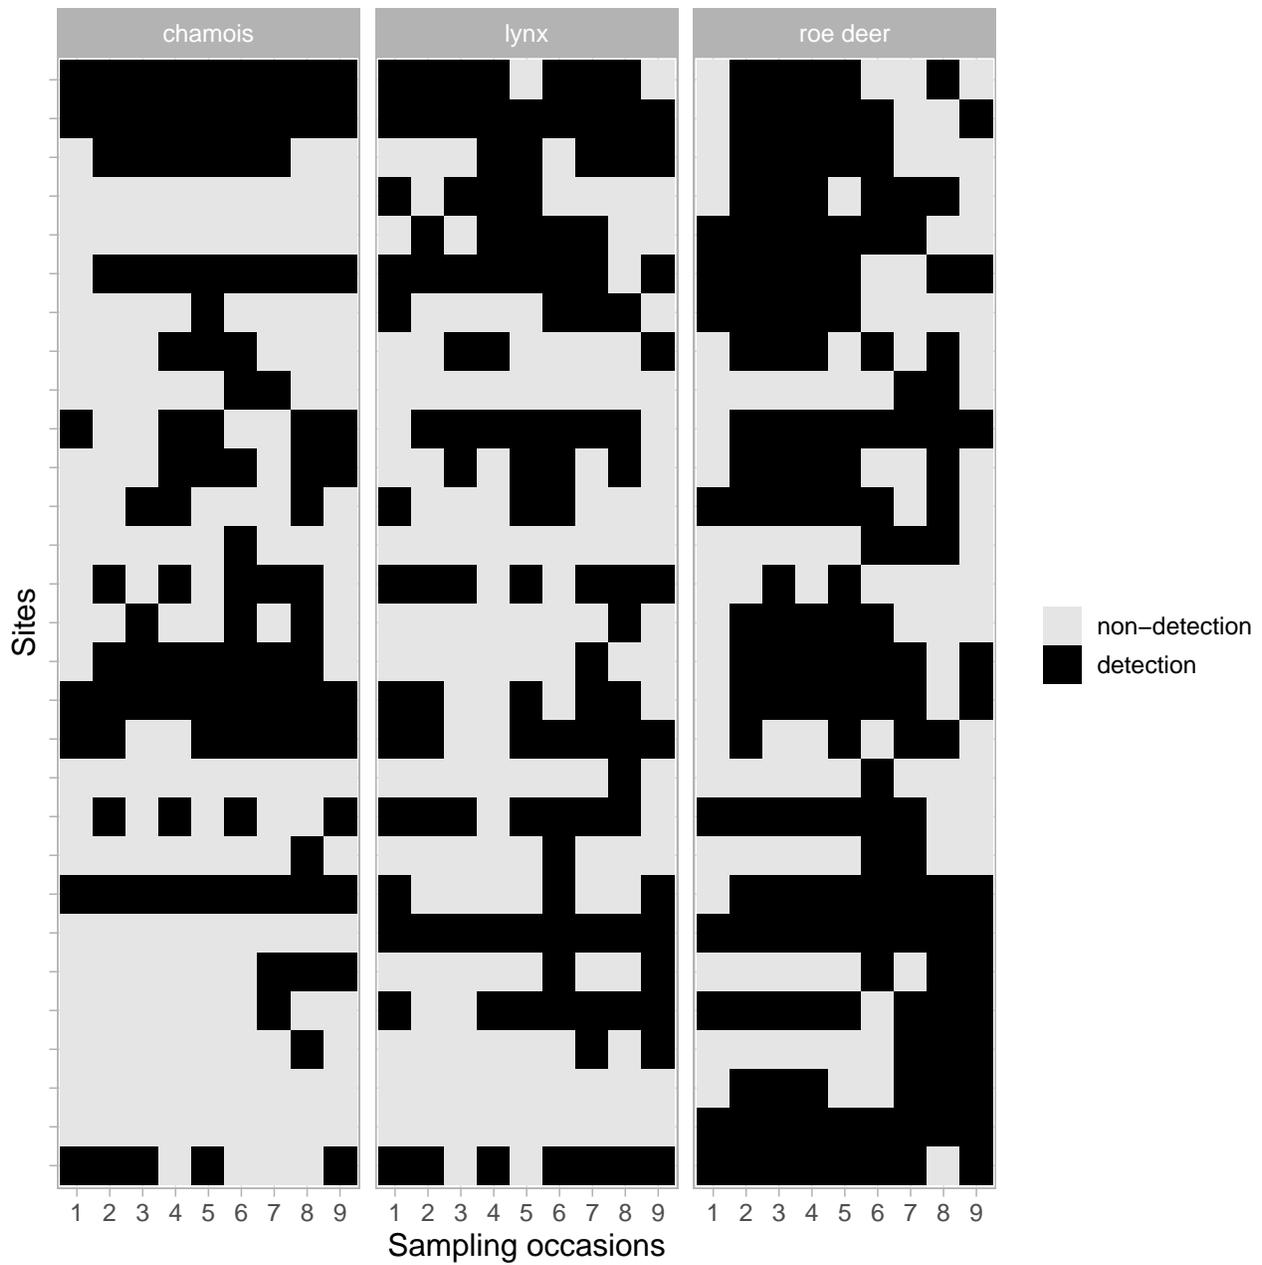


Figure 3: Detections (black) and non-detections (light grey) for each of the 3 species lynx, chamois and roe deer between March and November for all years pooled together. Sites are on the Y axis, while sampling occasions are on the X axis. Only data from the ground truth dataset are displayed.

(Dai, Ding, and Wahba 2013). Let us consider 2 species, species 1 and 2, then  $Z_i = (Z_{i1}, Z_{i2}) \sim$  multivariate Bernoulli( $\psi_{11}, \psi_{10}, \psi_{01}, \psi_{00}$ ) where  $\psi_{11}$  is the probability that a site is occupied by both species 1 and 2,  $\psi_{10}$  the probability that a site is occupied by species 1 but not 2,  $\psi_{01}$  the probability that a site is occupied by species 2 but not 1, and  $\psi_{00}$  the probability a site is occupied by none of them. Note that we considered species-specific only occupancy probabilities but these could be modeled as site-specific covariates. Marginal occupancy probabilities are obtained as  $\Pr(Z_{i1} = 1) = \psi_{11} + \psi_{10}$  and  $\Pr(Z_{i2} = 1) = \psi_{11} + \psi_{01}$ . With this model, we may also infer potential interactions by calculating conditional probabilities such as for example the probability of a site being occupied by species 2 conditional of species 1 with  $\Pr(Z_{i2} = 1|Z_{i1} = 1) = \frac{\psi_{11}}{\psi_{11} + \psi_{10}}$ .

Detection probabilities were indistinguishable whether we used the ground truth or the classified dataset, with  $p_{\text{lynx}} = 0.51(0.45, 0.58)$ ,  $p_{\text{roe deer}} = 0.63(0.57, 0.68)$  and  $p_{\text{chamois}} = 0.61(0.55, 0.67)$ .

We also found that occupancy probability estimates were similar whether we used the ground truth or the classified dataset (Figure 4). Roe deer was the prevalent species, but lynx and chamois were also occurring with high probability (Figure 4).

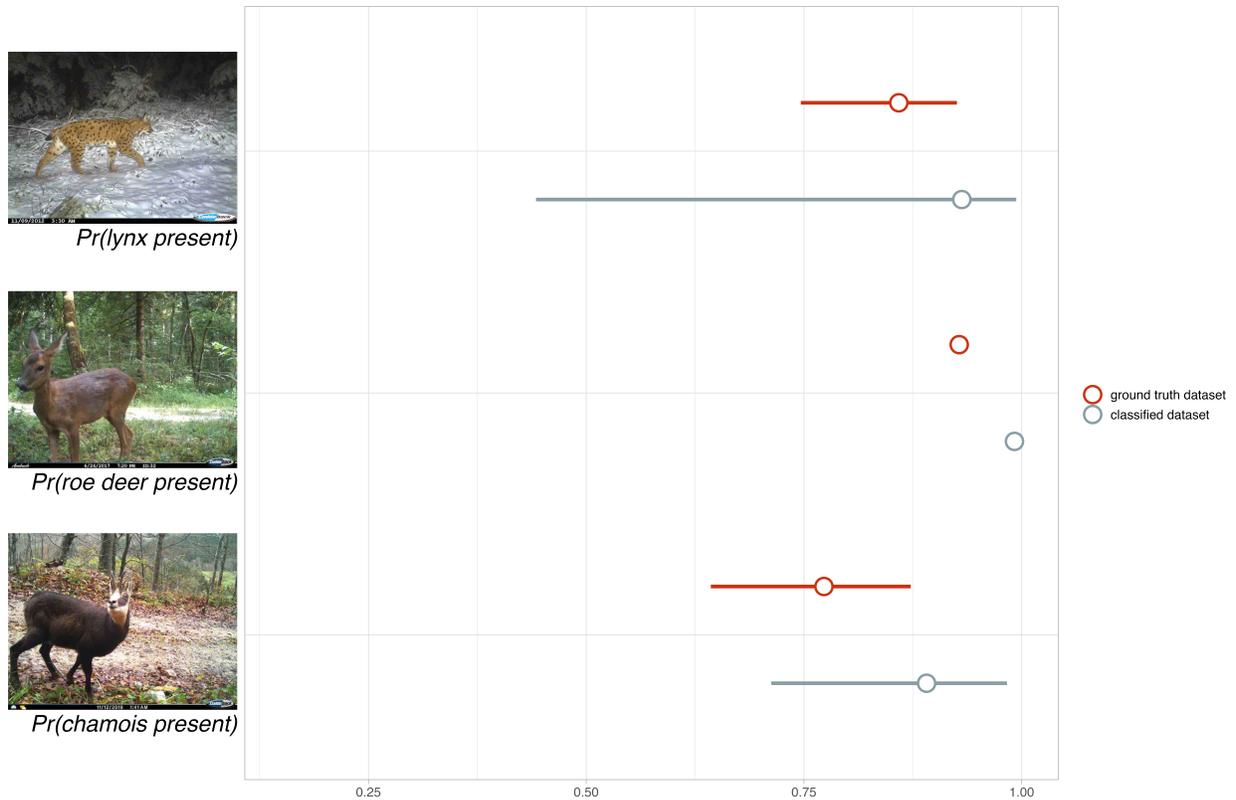


Figure 4: Marginal occupancy probabilities for all three species, lynx, roe deer and chamois). Parameter estimates are from a multispecies occupancy model using either the ground truth dataset (in red) or the classified dataset (in blue-grey).

Because marginal occupancy probabilities were high, probabilities of co-occurrence were also estimated high (Figure 5). Our results should be interpreted bearing in mind that co-occurrence is a necessary but not sufficient condition for actual interaction. When both preys were present, lynx was more present than when they were both absent (Figure 5). Lynx was more sensitive to the presence of roe deer than that of chamois (Figure 5).

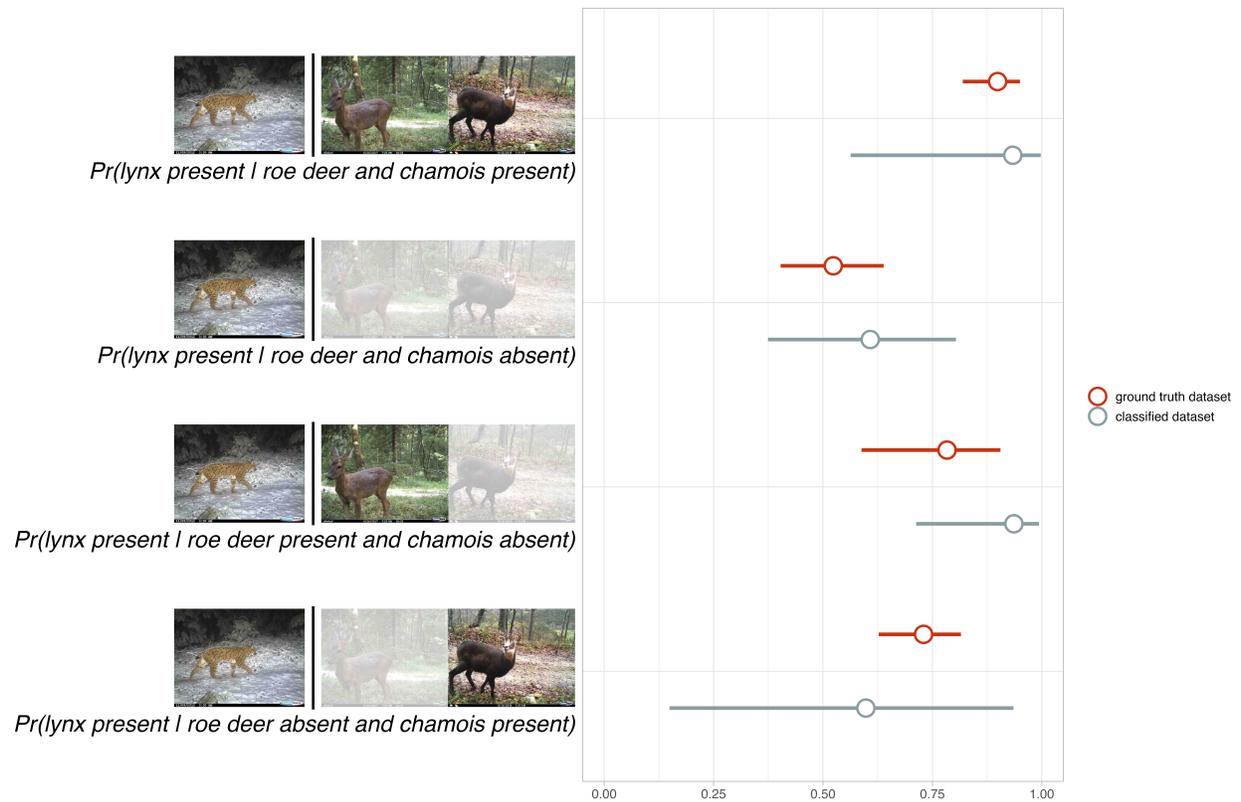


Figure 5: Lynx occupancy probability conditional on the presence or absence of its preys (roe deer and chamois). Parameter estimates are from a multispecies occupancy model using either the ground truth dataset (in red) or the classified dataset (in blue-grey).

## Discussion

In this paper, we aimed at illustrating a reproducible workflow for studying the structure of an animal community and species spatial co-occurrence (*why*) using images acquired from camera traps and automatically labelled with deep learning (*what*) which we analysed with statistical occupancy models accounting for imperfect species detection (*why*). Overall, we found that, even though model transferability could be improved, inference about the potential interactions between lynx and its preys was similar whether we analysed the ground truth data or classified data.

This result calls for further work on the trade-offs between time and resources allocated to train models with deep learning and our ability to correctly answer key ecological questions with camera-trap surveys. In other words, while a computer scientist might be keen on spending time training models to achieve top performances, an ecologist would rather rely on a model showing average performances and use this time to proceed with statistical analyses if, of course, errors in computer-annotated images do not make ecological inference flawed. The right balance may be found with collaborative projects in which scientists from artificial intelligence, statistics and ecology agree on a common objective, and identify research questions that can pick the interest of all parties.

Our demonstration remains however empirical, and we encourage others to try and replicate our results. We also see two avenues of research that could benefit the integration of deep learning and ecological statistics. First, a simulation study could be conducted to evaluate bias and precision in ecological parameter estimators with regard to errors in image annotation by computers. The outcome of this exercise could be, for example, guidelines informing on the confidence an investigator may place in ecological inference as a function of the amount of false negatives and false positives. Second, annotation errors could be accommodated directly in statistical models. For example, single-species occupancy models account for false negatives when a species is not detected by the camera at a site where it is present, as well as false positives when a species is detected at a site where it is not present due to species misidentification by the observer (Miller et al. 2011). Pending a careful distinction between ecological vs. computer-generated false negatives and false positives, error rates could be added to multispecies occupancy models (Chambert et al. 2018) and informed by recall and precision metrics obtained during model training (Tabak et al. 2020).

With regard to the case study, our results should be seen only as preliminary. First, we aim at quantifying the relative contribution of biotic (lynx predation on chamois and roe deer) and abiotic (habitat quality) processes to the composition and dynamic of this ecological community. Second, to benefit future camera trap studies of lynx in the Jura mountains, we plan to train a model again using manually annotated images from both the Jura and the Ain study sites. These perspectives are the object of ongoing work.

With the rapid advances in technologies for biodiversity monitoring (Lahoz-Monfort and Magrath 2021), the possibility of analysing large amounts of images makes deep learning appealing to ecologists. We hope that our proposal of a reproducible R workflow for deep learning and statistical ecology will encourage further studies in the integration of these disciplines, and contribute to the adoption of computer vision by ecologists.

## Appendix: Reproducible example of species identification on camera trap images with CPU

In this section, we go through a reproducible example of the entire deep learning workflow, including data preparation, model training, and automatic labeling of new images. We used a subsample of 467 images from the original dataset in the Jura county to allow the training of our model with CPU on a personal computer. We also used 14 images from the original dataset in the Ain county to illustrate prediction.

### Training and validation datasets

We first split the dataset of Jura images in two datasets, a dataset for training, and the other one for validation. We use the `exifr` package to extract metadata from images, get a list of images names and extract the species from these.

Table 4: Species considered, and number of images with these species in them.

Keywords	n
humain	143
vehicule	135
renard	58
sangliers	33
chasseur	17
chien	14
lynx	13
chevreuil	13
chamois	12
blaireaux	10
chat	8
lievre	4
fouine	1
cavalier	1

```

library(exifr)
pix_folder <- 'pix/pixJura/'
file_list <- list.files(path = pix_folder,
                        recursive = TRUE,
                        pattern = "*.JPG",
                        full.names = TRUE)

labels <-
  read_exif(file_list) %>%
  as_tibble() %>%
  unnest(Keywords, keep_empty = TRUE) %>% # keep_empty = TRUE keeps pix with no labels (empty pix)
  group_by(SourceFile) %>%
  slice_head() %>% # when several labels in a pix, keep first only
  ungroup() %>%
  mutate(Keywords = as_factor(Keywords)) %>%
  mutate(Keywords = fct_explicit_na(Keywords, "wo_tag")) %>% # when pix has no tag
  select(SourceFile, FileName, Keywords) %>%
  mutate(Keywords = fct_recode(Keywords,
                                "chat" = "chat forestier",
                                "lievre" = "lièvre",
                                "vehicule" = "véhicule",
                                "ni" = "Non identifié")) %>%
  filter(!(Keywords %in% c("ni", "wo_tag")))

```

Then we pick 80% of the images for training in each category, the rest being used for validation.

```

# training dataset
pix_train <- labels %>%
  select(SourceFile, FileName, Keywords) %>%
  group_by(Keywords) %>%
  filter(between(row_number(), 1, floor(n()*80/100))) # 80% per category

# validation dataset
pix_valid <- labels %>%
  group_by(Keywords) %>%
  filter(between(row_number(), floor(n()*80/100) + 1, n()))

```

Eventually, we store these images in two distinct directories named train and valid.

```
# create dir train/ and copy pix there, organised by categories
dir.create('pix/train') # create training directory
for (i in levels(fct_drop(pix_train$Keywords))) dir.create(paste0('pix/train/',i)) # create dir for lab
for (i in 1:nrow(pix_train)){
  file.copy(as.character(pix_train$SourceFile[i]),
            paste0('pix/train/', as.character(pix_train$Keywords[i]))) # copy pix in corresp dir
}
# create dir valid/ and copy pix there, organised by categories.
dir.create('pix/valid') # create validation dir
for (i in levels(fct_drop(pix_train$Keywords))) dir.create(paste0('pix/valid/',i)) # create dir for lab
for (i in 1:nrow(pix_valid)){
  file.copy(as.character(pix_valid$SourceFile[i]),
            paste0('pix/valid/', as.character(pix_valid$Keywords[i]))) # copy pix in corresp dir
}
# delete pictures in valid/ directory for which we did not train the model
to_be_deleted <- setdiff(levels(fct_drop(pix_valid$Keywords)), levels(fct_drop(pix_train$Keywords)))
if (!is_empty(to_be_deleted)) {
  for (i in 1:length(to_be_deleted)){
    unlink(paste0('pix/valid/', to_be_deleted[i]))
  }
}
```

What is the sample size of these two datasets?

```
bind_rows("training" = pix_train, "validation" = pix_valid, .id = "dataset") %>%
  group_by(dataset) %>%
  count(Keywords) %>%
  rename(category = Keywords) %>%
  kable(caption = "Sample size (n) for the training and validation datasets.") %>%
  kable_styling()
```

## Transfer learning

We proceed with transfer learning using images from the Jura county (or a subsample more exactly). We first load images and apply standard transformations to improve training (flip, rotate, zoom, rotate, light transform).

```
dls <- ImageDataLoaders_from_folder(
  path = "pix/",
  train = "train",
  valid = "valid",
  item_tfms = Resize(size = 460),
  bs = 10,
  batch_tfms = list(aug_transforms(size = 224,
                                   min_scale = 0.75), # transformation
                   Normalize_from_stats( imagenet_stats() )),
  num_workers = 0,
  ImageFile.LOAD_TRUNCATED_IMAGES = TRUE)
```

Then we get the model architecture. For the sake of illustration, we use a resnet18 here, but we used a resnet50 to get the full results presented in the main text.

```
learn <- cnn_learner(dls = dls,
                    arch = resnet18(),
```

Table 5: Sample size (n) for the training and validation datasets.

dataset	category	n
training	humain	114
training	vehicule	108
training	chamois	9
training	blaireaux	8
training	sangliers	26
training	renard	46
training	chasseur	13
training	lynx	10
training	chien	11
training	chat	6
training	chevreuil	10
training	lievre	3
validation	humain	29
validation	vehicule	27
validation	chamois	3
validation	blaireaux	2
validation	sangliers	7
validation	renard	12
validation	chasseur	4
validation	lynx	3
validation	chien	3
validation	fouine	1
validation	chat	2
validation	chevreuil	3
validation	lievre	1
validation	cavalier	1

```
metrics = list(accuracy, error_rate))
```

Now we are ready to train our model. Again, for the sake of illustration, we use only 2 epochs here, but used 20 epochs to get the full results presented in the main text. With all pictures and a resnet50, it took 75 minutes per epoch approximatively on a Mac with a 2.4Ghz processor and 64Go memory, and less than half an hour on a machine with GPU. On this reduced dataset, it took a bit more than a minute per epoch on the same Mac. Note that we save the model after each epoch for later use.

```
one_cycle <- learn %>%  
  fit_one_cycle(2, cbs = SaveModelCallback(every_epoch = TRUE,  
                                           fname = 'model'))
```

```
## epoch   train_loss  valid_loss  accuracy  error_rate  time  
## -----  -  
## 0        2.447199    0.850911   0.760417  0.239583   01:24  
## 1        1.736201    0.666055   0.781250  0.218750   01:25
```

```
one_cycle
```

```
##   epoch train_loss valid_loss  accuracy error_rate  
## 1     0   2.447199  0.8509111 0.7604167  0.2395833  
## 2     1   1.736201  0.6660554 0.7812500  0.2187500
```

We may dig a bit deeper in training performances by loading the best model, here `model_1.pth`, and display some metrics for each species.

```
learn$load("model_1")
```

```
## Sequential(  
##   (0): Sequential(  
##     (0): Conv2d(3, 64, kernel_size=(7, 7), stride=(2, 2), padding=(3, 3), bias=False)  
##     (1): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)  
##     (2): ReLU(inplace=True)  
##     (3): MaxPool2d(kernel_size=3, stride=2, padding=1, dilation=1, ceil_mode=False)  
##     (4): Sequential(  
##       (0): BasicBlock(  
##         (conv1): Conv2d(64, 64, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)  
##         (bn1): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)  
##         (relu): ReLU(inplace=True)  
##         (conv2): Conv2d(64, 64, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)  
##         (bn2): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)  
##       )  
##     (1): BasicBlock(  
##       (conv1): Conv2d(64, 64, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)  
##       (bn1): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)  
##       (relu): ReLU(inplace=True)  
##       (conv2): Conv2d(64, 64, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)  
##       (bn2): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)  
##     )  
##   )  
##   (5): Sequential(  
##     (0): BasicBlock(  
##       (conv1): Conv2d(64, 128, kernel_size=(3, 3), stride=(2, 2), padding=(1, 1), bias=False)  
##       (bn1): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)  
##       (relu): ReLU(inplace=True)  
##       (conv2): Conv2d(128, 128, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)  
##       (bn2): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
```

```

##         (downsample): Sequential(
##           (0): Conv2d(64, 128, kernel_size=(1, 1), stride=(2, 2), bias=False)
##           (1): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##         )
##       )
##     (1): BasicBlock(
##       (conv1): Conv2d(128, 128, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##       (bn1): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##       (relu): ReLU(inplace=True)
##       (conv2): Conv2d(128, 128, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##       (bn2): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     )
##   )
## (6): Sequential(
##   (0): BasicBlock(
##     (conv1): Conv2d(128, 256, kernel_size=(3, 3), stride=(2, 2), padding=(1, 1), bias=False)
##     (bn1): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     (relu): ReLU(inplace=True)
##     (conv2): Conv2d(256, 256, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##     (bn2): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     (downsample): Sequential(
##       (0): Conv2d(128, 256, kernel_size=(1, 1), stride=(2, 2), bias=False)
##       (1): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     )
##   )
##   (1): BasicBlock(
##     (conv1): Conv2d(256, 256, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##     (bn1): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     (relu): ReLU(inplace=True)
##     (conv2): Conv2d(256, 256, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##     (bn2): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##   )
## )
## (7): Sequential(
##   (0): BasicBlock(
##     (conv1): Conv2d(256, 512, kernel_size=(3, 3), stride=(2, 2), padding=(1, 1), bias=False)
##     (bn1): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     (relu): ReLU(inplace=True)
##     (conv2): Conv2d(512, 512, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##     (bn2): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     (downsample): Sequential(
##       (0): Conv2d(256, 512, kernel_size=(1, 1), stride=(2, 2), bias=False)
##       (1): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     )
##   )
##   (1): BasicBlock(
##     (conv1): Conv2d(512, 512, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##     (bn1): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##     (relu): ReLU(inplace=True)
##     (conv2): Conv2d(512, 512, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1), bias=False)
##     (bn2): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
##   )
## )
## )

```

```
## (1): Sequential(
## (0): AdaptiveConcatPool2d(
## (ap): AdaptiveAvgPool2d(output_size=1)
## (mp): AdaptiveMaxPool2d(output_size=1)
## )
## (1): Flatten(full=False)
## (2): BatchNorm1d(1024, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
## (3): Dropout(p=0.25, inplace=False)
## (4): Linear(in_features=1024, out_features=512, bias=False)
## (5): ReLU(inplace=True)
## (6): BatchNorm1d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
## (7): Dropout(p=0.5, inplace=False)
## (8): Linear(in_features=512, out_features=12, bias=False)
## )
## )

interp <- ClassificationInterpretation_from_learner(learn)
interp$print_classification_report()
```

We may extract the categories that get the most confused.

```
interp %>% most_confused()
```

```
##           V1           V2 V3
## 1    humain  vehicule  4
## 2   chasseur   humain  3
## 3  blaireaux   renard  1
## 4  blaireaux sangliers  1
## 5         chat   renard  1
## 6         chat sangliers  1
## 7  chevreuil   chien  1
## 8  chevreuil   renard  1
## 9         chien  chamois  1
## 10         chien sangliers  1
## 11         humain chasseur  1
## 12         lievre   renard  1
## 13         renard blaireaux  1
## 14         renard sangliers  1
## 15 sangliers   renard  1
## 16 vehicule   humain  1
```

## Transferability

In this section, we show how to use our freshly trained model to label images that were taken in another study site in the Ain county, and not used to train our model. First, we get the path to the images.

```
fls <- list.files(path = "pix/pixAin",
                 full.names = TRUE,
                 recursive = TRUE)
```

Then we carry out prediction, and compare to the truth.

```
predicted <- character(3)
categories <- interp$vocab %>%
  str_replace_all("[[:punct:]]", " ") %>%
  str_trim() %>%
  str_split(" ") %>%
```

Table 6: Comparison of the predictions vs. ground truth.

truth	prediction
lynx	chevreuil
roe deer	chamois
wild boar	sangliers

```

unlist()
for (i in 1:length(fls)){
  result <- learn %>% predict(fls[i]) # make prediction
  result[[3]] %>%
    str_extract("\\d+") %>%
    as.integer() -> index # extract relevant info
  predicted[i] <- categories[index + 1] # match it with categories
}
data.frame(truth = c("lynx", "roe deer", "wild boar"),
           prediction = predicted) %>%
  kable(caption = "Comparison of the predictions vs. ground truth.") %>%
  kable_styling()

```

## Session information

```

## R version 4.1.0 (2021-05-18)
## Platform: x86_64-apple-darwin17.0 (64-bit)
## Running under: macOS Catalina 10.15.7
##
## Matrix products: default
## BLAS: /Library/Frameworks/R.framework/Versions/4.1/Resources/lib/libRblas.dylib
## LAPACK: /Library/Frameworks/R.framework/Versions/4.1/Resources/lib/libRlapack.dylib
##
## locale:
## [1] fr_FR.UTF-8/fr_FR.UTF-8/fr_FR.UTF-8/C/fr_FR.UTF-8/fr_FR.UTF-8
##
## attached base packages:
## [1] stats graphics grDevices utils datasets methods base
##
## other attached packages:
## [1] exifr_0.3.2 unmarked_1.1.1.9006 lattice_0.20-44
## [4] janitor_2.1.0 highcharter_0.8.2 fastai_2.0.9
## [7] ggtext_0.1.1 wesanderson_0.3.6 kableExtra_1.3.4
## [10] stringi_1.7.3 lubridate_1.7.10 cowplot_1.1.1
## [13] sf_1.0-2 forcats_0.5.1 stringr_1.4.0
## [16] dplyr_1.0.7 purrr_0.3.4.9000 readr_2.0.0
## [19] tidyr_1.1.3 tibble_3.1.3 ggplot2_3.3.5
## [22] tidyverse_1.3.1
##
## loaded via a namespace (and not attached):
## [1] minqa_1.2.4 colorspace_2.0-2 ggsignif_0.6.2
## [4] rio_0.5.27 ellipsis_0.3.2 class_7.3-19
## [7] snakecase_0.11.0 markdown_1.1 fs_1.5.0
## [10] gridtext_0.1.4 rstudioapi_0.13 proxy_0.4-26

```

```

## [13] ggpubr_0.4.0      farver_2.1.0      bit64_4.0.5
## [16] fansi_0.5.0       xml2_1.3.2       codetools_0.2-18
## [19] splines_4.1.0     knitr_1.33       rlist_0.4.6.1
## [22] jsonlite_1.7.2    nloptr_1.2.2.2   broom_0.7.9
## [25] dbplyr_2.1.1      png_0.1-7        compiler_4.1.0
## [28] httr_1.4.2        backports_1.2.1  assertthat_0.2.1
## [31] Matrix_1.3-3      cli_3.0.1        htmltools_0.5.1.1
## [34] tools_4.1.0       igraph_1.2.6     gtable_0.3.0
## [37] glue_1.4.2        rappdirs_0.3.3   Rcpp_1.0.7
## [40] carData_3.0-4     cellranger_1.1.0 raster_3.4-13
## [43] vctrs_0.3.8       nlme_3.1-152     svglite_2.0.0
## [46] xfun_0.25         ps_1.6.0         openxlsx_4.2.4
## [49] lme4_1.1-27.1     rvest_1.0.1      lifecycle_1.0.0
## [52] rstatix_0.7.0     MASS_7.3-54      zoo_1.8-9
## [55] scales_1.1.1      vroom_1.5.4     ragg_1.1.3
## [58] hms_1.1.0         parallel_4.1.0   yaml_2.2.1
## [61] quantmod_0.4.18  curl_4.3.2       reticulate_1.20
## [64] pbapply_1.4-3     highr_0.9        e1071_1.7-8
## [67] TTR_0.24.2        zip_2.2.0        boot_1.3-28
## [70] rlang_0.4.11      pkgconfig_2.0.3  systemfonts_1.0.2
## [73] evaluate_0.14     labeling_0.4.2   htmlwidgets_1.5.3
## [76] bit_4.0.4         tidyselect_1.1.1 processx_3.5.2
## [79] plyr_1.8.6        magrittr_2.0.1   R6_2.5.0
## [82] generics_0.1.0    DBI_1.1.1        foreign_0.8-81
## [85] pillar_1.6.2      haven_2.4.3      withr_2.4.2
## [88] units_0.7-2       xts_0.12.1       abind_1.4-5
## [91] sp_1.4-5          car_3.0-11       modelr_0.1.8
## [94] crayon_1.4.1      KernSmooth_2.23-20 utf8_1.2.2
## [97] tzdb_0.1.2        rmarkdown_2.10   jpeg_0.1-9
## [100] grid_4.1.0        readxl_1.3.1     data.table_1.14.0
## [103] callr_3.7.0       reprex_2.0.1     digest_0.6.27
## [106] classInt_0.4-3    webshot_0.5.2    textshaping_0.3.5
## [109] munsell_0.5.0     viridisLite_0.4.0

```

## Acknowledgments

We warmly thank Mathieu Massaviol, Remy Dernas and Khalid Belkhir for their help in using GPU machines on the Montpellier Bioinformatics Biodiversity platform, Julien Renoult for helpful discussions, Delphine Dinouart and Chloé Quillard for their precious help in manually tagging the images, and Vincent Miele for having inspired this work, and his help and support along the way. We also thank the staff of the Federations of Hunters from the Jura and Ain counties, hunters who helped to find locations for camera traps and volunteers who contributed in collecting data. Last, we thank Auvergne-Rhône-Alpes Region, Ain and Jura departmental Councils, The French National Federation of Hunters, French Environmental Ministry based in Auvergne-Rhône-Alpes and Bourgogne Franche-Comté Region and the French Office for Biodiversity for funding the Lynx Predator Prey Program. This work was also partly funded by the French National Research Agency (grant ANR-16-CE02-0007).

## References

Allaire, JJ, Kevin Ushey, Yuan Tang, and Dirk Eddelbuettel. 2017. *Reticulate: R Interface to Python*. <https://github.com/rstudio/reticulate>.

- Baraniuk, Richard, David Donoho, and Matan Gavish. 2020. “The Science of Deep Learning.” *Proceedings of the National Academy of Sciences* 117 (48): 30029–32. <https://doi.org/10.1073/pnas.2020596117>.
- Beery, Sara, Grant van Horn, and Pietro Perona. 2018. “Recognition in Terra Incognita.” *arXiv:1807.04975 [Cs, q-Bio]*, July. <http://arxiv.org/abs/1807.04975>.
- Breitenmoser, Urs. 1998. “Large Predators in the Alps: The Fall and Rise of Man’s Competitors.” *Biological Conservation, Conservation Biology and Biodiversity Strategies*, 83 (3): 279–89. [https://doi.org/10.1016/S0006-3207\(97\)00084-0](https://doi.org/10.1016/S0006-3207(97)00084-0).
- Chambert, Thierry, Evan H. Campbell Grant, David A. W. Miller, James D. Nichols, Kevin P. Mulder, and Adrienne B. Brand. 2018. “Two-Species Occupancy Modelling Accounting for Species Misidentification and Non-Detection.” *Methods in Ecology and Evolution* 9 (6): 1468–77. <https://doi.org/https://doi.org/10.1111/2041-210X.12985>.
- Christin, Sylvain, Éric Hervet, and Nicolas Lecomte. 2019. “Applications for Deep Learning in Ecology.” Edited by Hao Ye. *Methods in Ecology and Evolution* 10 (10): 1632–44. <https://doi.org/10.1111/2041-210X.13256>.
- Clipp, Hannah L., Amber L. Evans, Brin E. Kessinger, K. Kellner, and Christopher T. Rota. 2021. “A Penalized Likelihood for Multi-Species Occupancy Models Improves Predictions of Species Interactions.” *Ecology*.
- Dai, Bin, Shilin Ding, and Grace Wahba. 2013. “Multivariate Bernoulli Distribution.” *Bernoulli* 19 (4). <https://doi.org/10.3150/12-BEJSP10>.
- Duggan, Matthew T., Melissa F. Groleau, Ethan P. Shealy, Lillian S. Self, Taylor E. Utter, Matthew M. Waller, Bryan C. Hall, Chris G. Stone, Layne L. Anderson, and Timothy A. Mousseau. 2021. “An Approach to Rapid Processing of Camera Trap Images with Minimal Human Input.” *Ecology and Evolution*. <https://doi.org/https://doi.org/10.1002/ece3.7970>.
- Fiske, Ian, and Richard Chandler. 2011. “unmarked: An R Package for Fitting Hierarchical Models of Wildlife Occurrence and Abundance.” *Journal of Statistical Software* 43 (10): 1–23. <https://www.jstatsoft.org/v43/i10/>.
- Gimenez, Olivier, Stephen T. Buckland, Byron J. T. Morgan, Nicolas Bez, Sophie Bertrand, Rémi Choquet, Stéphane Dray, et al. 2014. “Statistical Ecology Comes of Age.” *Biology Letters* 10 (12): 20140698. <https://doi.org/10.1098/rsbl.2014.0698>.
- He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. “Deep Residual Learning for Image Recognition.” In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–78. <https://doi.org/10.1109/CVPR.2016.90>.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2012. “ImageNet Classification with Deep Convolutional Neural Networks.” In *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, 1097–1105. Curran Associates, Inc.
- Lahoz-Monfort, José J, and Michael J L Magrath. 2021. “A Comprehensive Overview of Technologies for Species and Habitat Monitoring and Conservation.” *BioScience*. <https://doi.org/10.1093/biosci/biab073>.
- Lai, Jiangshan, Christopher J. Lortie, Robert A. Muenchen, Jian Yang, and Keping Ma. 2019. “Evaluating the Popularity of R in Ecology.” *Ecosphere* 10 (1). <https://doi.org/10.1002/ecs2.2567>.
- Lamba, Aakash, Phillip Cassey, Ramesh Raja Segaran, and Lian Pin Koh. 2019. “Deep Learning for Environmental Conservation.” *Current Biology* 29 (19): R977–82. <https://doi.org/10.1016/j.cub.2019.08.016>.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. “Deep Learning.” *Nature* 521 (7553): 436–44. <https://doi.org/10.1038/nature14539>.
- Miller, David A., James D. Nichols, Brett T. McClintock, Evan H. Campbell Grant, Larissa L. Bailey, and Linda A. Weir. 2011. “Improving Occupancy Estimation When Two Types of Observational Error

- Occur: Non-Detection and Species Misidentification.” *Ecology* 92 (7): 1422–28. <https://doi.org/https://doi.org/10.1890/10-1396.1>.
- Molinari-Jobin, Anja, Fridolin Zimmermann, Andreas Ryser, Christine Breitenmoser-Würsten, Simon Capt, Urs Breitenmoser, Paolo Molinari, Heinrich Haller, and Roman Eyholzer. 2007. “Variation in Diet, Prey Selectivity and Home-Range Size of Eurasian Lynx *Lynx Lynx* in Switzerland.” *Wildlife Biology* 13 (4): 393–405. [https://doi.org/10.2981/0909-6396\(2007\)13%5B393:VIDPSA%5D2.0.CO;2](https://doi.org/10.2981/0909-6396(2007)13%5B393:VIDPSA%5D2.0.CO;2).
- Norouzzadeh, Mohammad Sadegh, Dan Morris, Sara Beery, Neel Joshi, Nebojsa Jojic, and Jeff Clune. 2021. “A Deep Active Learning System for Species Identification and Counting in Camera Trap Images.” Edited by Matthew Schofield. *Methods in Ecology and Evolution* 12 (1): 150–61. <https://doi.org/10.1111/2041-210X.13504>.
- Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, et al. 2019. “PyTorch: An Imperative Style, High-Performance Deep Learning Library.” In *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, 8024–35. Curran Associates, Inc. <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Rota, Christopher T., Marco A. R. Ferreira, Roland W. Kays, Tavis D. Forrester, Elizabeth L. Kalies, William J. McShea, Arielle W. Parsons, and Joshua J. Millspaugh. 2016. “A Multispecies Occupancy Model for Two or More Interacting Species.” *Methods in Ecology and Evolution* 7 (10): 1164–73. <https://doi.org/https://doi.org/10.1111/2041-210X.12587>.
- Shao, Ling, Fan Zhu, and Xuelong Li. 2015. “Transfer Learning for Visual Categorization: A Survey.” *IEEE Transactions on Neural Networks and Learning Systems* 26 (5): 1019–34. <https://doi.org/10.1109/TNNLS.2014.2330900>.
- Shorten, Connor, and Taghi M. Khoshgoftaar. 2019. “A Survey on Image Data Augmentation for Deep Learning.” *Journal of Big Data* 6 (1): 60. <https://doi.org/10.1186/s40537-019-0197-0>.
- Sutherland, William J., Robert P. Freckleton, H. Charles J. Godfray, Steven R. Beissinger, Tim Benton, Duncan D. Cameron, Yohay Carmel, et al. 2013. “Identification of 100 Fundamental Ecological Questions.” Edited by David Gibson. *Journal of Ecology* 101 (1): 58–67. <https://doi.org/10.1111/1365-2745.12025>.
- Tabak, Michael A., Mohammad S. Norouzzadeh, David W. Wolfson, Erica J. Newton, Raoul K. Boughton, Jacob S. Ivan, Eric A. Odell, et al. 2020. “Improving the Accessibility and Transferability of Machine Learning Algorithms for Identification of Animals in Camera Trap Images: Mlwic2.” *Ecology and Evolution* 10 (19): 10374–83. <https://doi.org/10.1002/ece3.6692>.
- Tabak, Michael A., Mohammad S. Norouzzadeh, David W. Wolfson, Steven J. Sweeney, Kurt C. Vercauteren, Nathan P. Snow, Joseph M. Halseth, et al. 2019. “Machine Learning to Classify Animal Species in Camera Trap Images: Applications in Ecology.” Edited by Theoni Photopoulou. *Methods in Ecology and Evolution* 10 (4): 585–90. <https://doi.org/10.1111/2041-210X.13120>.
- Vandel, Jean-Michel, and Philippe Stahl. 2005. “Distribution Trend of the Eurasian Lynx *Lynx Lynx* Populations in France.” *Mammalia* 69 (2). <https://doi.org/10.1515/mamm.2005.013>.
- Voulodimos, Athanasios, Nikolaos Doulamis, Anastasios Doulamis, and Eftychios Protopapadakis. 2018. “Deep Learning for Computer Vision: A Brief Review.” Edited by Diego Andina. *Computational Intelligence and Neuroscience* 2018 (February): 7068349. <https://doi.org/10.1155/2018/7068349>.
- Wearn, Oliver R., Robin Freeman, and David M. P. Jacoby. 2019. “Responsible AI for Conservation.” *Nature Machine Intelligence* 1 (2): 72–73. <https://doi.org/10.1038/s42256-019-0022-7>.
- Weinstein, Ben G. 2018. “A Computer Vision for Animal Ecology.” Edited by Laura Prugh. *Journal of Animal Ecology* 87 (3): 533–45. <https://doi.org/10.1111/1365-2656.12780>.
- Willi, Marco, Ross T. Pitman, Anabelle W. Cardoso, Christina Locke, Alexandra Swanson, Amy Boyer, Marten Veldthuis, and Lucy Fortson. 2019. “Identifying Animal Species in Camera Trap Images Using

Deep Learning and Citizen Science.” Edited by Oscar Gaggiotti. *Methods in Ecology and Evolution* 10 (1): 80–91. <https://doi.org/10.1111/2041-210X.13099>.

Yosinski, Jason, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. “How Transferable Are Features in Deep Neural Networks?” In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, 3320–28. NIPS’14. Cambridge, MA, USA: MIT Press.

Zimmermann, Fridolin, Christine Breitenmoser-Würsten, Anja Molinari-Jobin, and Urs Breitenmoser. 2013. “Optimizing the Size of the Area Surveyed for Monitoring a Eurasian Lynx (*Lynx Lynx*) Population in the Swiss Alps by Means of Photographic Capture-Recapture.” *Integrative Zoology* 8 (3): 232–43. <https://doi.org/10.1111/1749-4877.12017>.