

ORIGINAL PAPER

SnakeFace: a transfer learning based app for snake classification

Jorge Guerra Pires ¹ and Luiz Henrique Dias Braga ²

¹Founder at IdeaCoding Lab / JovemPesquisador.com, Brazil, ²Laboratório de Zoologia dos Vertebrados, Departamento de Biodiversidade, Evolução e Meio Ambiente/ICEB, Universidade Federal de Ouro Preto, Campus Morro do Cruzeiro, s/n, Bauxita, Ouro Preto 35400-000, MG, Brazil

*jorgeguerrapires@yahoo.com.br; luizbraga.bio@gmail.com

Received: 2023-07-06. Revised: 2023-10-30. Accepted: 2023-11-16.

Abstract

Introduction: deep learning emerged in 2012 as one of the most important machine learning technologies, reducing image identification error from 25% to 5%. This article has two goals: 1) to demonstrate to the general public the ease of building state-of-the-art machine learning models without coding expertise; 2) to present a basic model adaptable to any biological image identification, such as species identification. Method: We present three test-of-concept models that showcase distinct perspectives of the app. The models aim at separating images into classes such as genus, species, and subspecies, and the input images can be easily adapted for different cases. We have applied deep learning and transfer learning using Teachable Machine. Results: Our basic models demonstrate high accuracy in identifying different species based on images, highlighting the potential for this method to be applied in biology. Discussions: the presented models showcase the ease of using machine learning nowadays for image identification. Furthermore, the adaptability of this method to various species and genres emphasizes its importance in the biological fields, as root for inspiring collaborations with computer science. On our case, future collaborations could lead to increasingly accurate and efficient models in this arena using well-curated datasets.

Keywords: Biology; bioinformatics; deep learning; Snakes; tensorflow; JavaScript; transfer learning

Resumo

Introdução: O aprendizado profundo surgiu em 2012 como uma das tecnologias mais importantes de aprendizado de máquina, reduzindo o erro de identificação de imagens de 25% para 5%. Este artigo tem dois objetivos: 1) demonstrar ao público em geral a facilidade de construir modelos de aprendizado de máquina de última geração sem expertise em programação; 2) apresentar um modelo básico adaptável a qualquer identificação de imagens biológicas, como identificação de espécies. Método: Apresentamos três modelos conceituais de teste que mostram perspectivas distintas do aplicativo. Os modelos visam separar imagens em classes como gênero, espécie e subespécie, e as imagens de entrada podem ser facilmente adaptadas para diferentes casos. Aplicamos aprendizado profundo e transferência de aprendizado usando o *Teachable Machine*. Resultados: Nossos modelos básicos demonstram alta precisão na identificação de diferentes espécies com base em imagens, destacando o potencial deste método para ser aplicado em biologia. Discussões: os modelos apresentados mostram a facilidade de usar aprendizado de máquina atualmente para identificação de imagens. Além disso, a adaptabilidade desse método a várias espécies e gêneros enfatiza sua importância nas áreas biológicas, como base para colaborações inspiradoras com ciência da computação. No nosso caso, colaborações futuras poderiam levar a modelos cada vez mais precisos e eficientes nessa área usando conjuntos de dados bem selecionados.

Palavras-Chave: Aprendizado profundo; bioinformática; biologia; JavaScript; Serpentes; TensorFlow; transferência de aprendizado.

1 Introduction

Around 2012, deep learning came to the surface, and became one of the most important machine learning technologies. The image identification error went from about 25% to just about 5%. Those values can be higher, or at least close to human error rates on image classification. This revolution also went to other areas. Google Translator, chatGPT, and other tools, they all use deep learning. China announced heavy investments in deep learning researches, after they were beaten on their own game Go using deep learning ([FRONTLINE, 2020](#)).

Different from human intelligence, machine learning is stored and easily transferable. When Einstein died, so died his genius. When we train a model, you can either give access to the model (e.g. chatGPT) to the public, called pretrained models, or actually borrow part of the model, called transfer learning. These approaches make machine learning easily accessible to anyone, from anywhere. One of the selling points of chatGPT is making large language models accessible, models that not even organizations could afford on their own. Said the CEO of openAI ([Altman, 2023](#)): the cost of intelligence will drop drastically in the upcoming years.

Nowadays, we have a common consensus that AI should be accessible to anyone; not just experts, not just big techs. Generally, they are released as APIs (Application Programming Interface), or public libraries. In a close past, those models were just for experts, and were very hard to work with, very focused on all aspects, including programming languages. Good models were generally locked up on highly-cost proprietary software. With APIs, applications can talk to each and exchange information and abilities, or even, different AIs, and share easily what they have learnt, what they do best. chatGPT has been constantly integrated with other tools using APIs, even, other Artificial Intelligence (AIs) models. It makes those models exponentially more powerful, and more accessible.

1.1 Our goal

This article aims at two goals, at its core: 1) showing to the general public how straightforward it is nowadays to build machine learning models with state of the art techniques, no code required, no expertise required; 2) present a basic model that can be adapted to any similar case in biology (i.e., species identification based on images, as long as the identification can be done from images, the model should work). Also, we want to draw attention to possible collaborations.

We also propose a big model, which we hope will be able to support people on identifying snakes. Our mindset is similar to chatGPT developers: launch it as soon as possible, and get feedback, even if negative. Our hope is that we could grow together with the users, instead of following the traditional path of just launching when the app is perfect. We have launched a simple demo, a test of concept, and share our findings on this article. As we are going to show, the app misclassifies when the snakes are similar, and we have a couple of thoughts on how to improve the predictions. To be fair, where the model misclassifies, non-experts also may misclassify (e.g., false

coral snakes vs. true ones).

1.2 Basic concepts

Our main bottleneck is a well-curated dataset of snakes with species and subspecies, we are using Google Image which can introduce errors of classification from erroneous annotations. Even though the machine learning training process is set generally to ignore noise on dataset (i.e., mistagged snakes), we would like to avoid those noises when possible. On machine learning community, we call this process of associating classes to images, annotation, the process of annotating images for teaching an algorithm to do it on its own.

Annotation can be a very big problem since it may take time, and time from experts. That is where the quotation from the CEO from openAI Sam Altman comes handy, once those models are trained, they can perform the work of experts, at zero or low costs; chatGPT costs one cent of dollar per 1.000 words, no human can beat that. Our dataset was partially curated by a biologist, but enlarged with images from Google Images. Since we are using transfer learning, about 30 images for each class was enough; image models are trained generally on millions of images for being good enough.

A nice move we did was using transfer learning: it makes the annotation less heavy since we need a small set of samples for each class. A machine learning may require millions of images for actually achieving interesting results. MobileNetV1, as an example, was trained on more than 1 million images.

Transfer learning was an ingenious way found by computer scientists to solve two problems: i) making it less heavy to train those models; ii) making it available already trained models, but, with learning abilities.

On this approach, we have a base model, and a trainable model. The base model is trained heavily, and made available online for users, which use this base model as feature model (see [Section 2.7](#)). A trainable model will use those features for training a local model. Our feature model has about 1.200 features, which is sampled from images. Models made available on [TensorFlow Hub](#) can have more than 1600 features, which is a measure on how many details they can learn from images. TensorFlow Hub is a collection of pretrained models made available for users, no charge, feature models are one type of models made available for transfer learning on this Google platform. Their goal is making transfer learning accessible, and easy to use, especially on the browser (i.e., JavaScript). TensorFlow.js is also part of this pursuit of making machine learning widely available, it was created by a group at Google, the same creators of TensorFlow in python (they are similar, and exchangeable).

As one example on how this tool can be powerful, we have [Pires \(2018\)](#) reported an example where researchers applied a general-trained based model for teeth segmentation, making it possible to achieve excellent results from a small number of images. We are using the same approach, except they had to build the model from scratch, requiring a couple of PhD students and masters in computer science as so they could make the transfer learning. On its basics, transfer is not

straightforward to implement.

We are going to use TensorFlow.js (Rivera, 2020; Shanjing Cai et al., 2020; Laborde, 2021). TensorFlow.js was developed by Google, based on TensorFlow in python, also by Google. Their goal was making machine learning more accessible, also for web developers. Moreover, taking advantage from the most used programming environment: the browser; and the most used programming language, JavaScript (Stack Overflow, 2023). The nice point about JavaScript is that most modern libraries and frameworks for web applications are somehow rooted in JavaScript, making it easy any extension. Our app was developed in Angular, also by Google, coded in Typescript from Microsoft, a typed version of JavaScript. Angular can be useful for scientific computation, as it was pointed out here (Pires et al., 2021).

1.2.1 On the importance of having an app for identifying snakes even by non-experts

In 2022, 27515 accidents with snakes were notified in Brazil, with death rate of about 0,40% (Sistema de Informação de Agravos e de Notificações – SINAN, 2023). Some numbers given by the organization: about 69% are attributed to *Bothrops* genus (simulations showed a good identification of this genus by our model since they are very easy to define on images); 1,19% to *Micrurus* (our model is not very good at separating from the fake ones; the goal should be avoiding misclassifying the true one as false coral snake); 16,6% to another snake genus and 13,4% non-identified snakes.

According to the healthcare agency from Paraná (Brazil) (Secretaria de Saúde, 2022), in Brazil are notified annually about 20.0000 accidents with snakes, with death rate of about 0,43%. Some numbers given by the organization: 1) about 75% are attributed to *bothrops* genus (simulations showed a good identification of this genus by our model since they are very easy to define on images); 2) *Micrurus* (0,5% our model is not very good at separating from the fake ones; the goal should be avoiding misclassifying the true one as false coral snake). Moreover, 16,6% to another snake genus and 13,4% non-identified snakes.

In this fact, services to help on identification of snakes is important and could minimize risks of accidents or misidentification during the treatment, since some non-venomous snakes could be morphologically similar with venomous snakes.

The snake genus *Micrurus*, commonly known as coral snake, is a well-known snake by the Brazilian population for its color pattern: mostly of the species in the genus presents red, black, and white (or yellow) rings (Roze, 1996). This type of color pattern is not by chance: its biological evolution for advertise their dangerous called as “warning coloration” (i.e., aposematism, Poulton (1890)). In contrast, certain non-harmful species (mimics) imitates the visual, acoustic, and chemical characteristics of toxic species (models). This phenomenon is called Batesian mimicry (Ruxton et al. (2004); Hossie and Sherratt (2013); Davis Rabosky et al. (2016)). Within this context, species belonging to the *Micrurus* genus are often associated with being mimic models for non-venomous or “semi-venomous”

snakes, commonly known as false coral snakes (Wickler, 1968; Greene and McDiarmid, 1981; Campbell and Lamar, 1989; Savage and Slowinski, 1992). Some of these species mimic various *Micrurus* species, showcasing different color patterns in according to their distribution and of the mimic model (e.g., *Atracus* sp., Almeida et al. (2022)).

1.2.2 A brief look at the literature

Even though we were the first to explore transfer learning for snake classification using JavaScript, the application of machine learning to snake classification, that we explore herein, has already been explored also by other researchers. A quick search on the internet, focusing on technical works, we can find the following researches. This topic has the potential to support on avoiding accidents with snakes, and even protecting the non-poisonous snakes.

Rajabizadeh and Rezghi (2021) did a work very similar to ours, and compared several algorithms; they have used as based model MobileNetV2. They also compared different methods: one of them is using Principal Component Analysis (PCA) for supporting the learning process. The feature model is a kind of PCA, when one reduces an image to a number of features, it is a sort of PCA. Their best results was in the same direction as ours: transfer learning based.

Kalinathan et al. (2021) was able to classify 772 classes of snake species using the architecture ResNeXt50-V2, which is an impressive result. In our case, we do not know our architecture since this is not available on Teachable Machine, but we know it has about 1200 features by analyzing their pretrained model locally. Laborde (2021) suggests that the final model is “a transfer learning version of MobileNet, PoseNet, or some other practical model that fits your needs.”

Ahmed et al. (2023) tested several base models on snake classification, and MobieNet was one of the top algorithms, the one we are using as base model for transfer learning.

Yang and Sinnott (2021) created a model for 11 snake breeds from Australia, and also made it available as iOS application; for our case, we are betting on a Reactive Web Design (RWD), no need to keep two or more codes, the app supposes to chance accordingly. Their lower number on classes seems to be due to updated in machine learning, much has changed in four years since the publication.

Durso et al. (2021) created a model for 45 species of snakes, they tried several models.

Santos et al. (2021) did something very similar to us. Their app Snake Classifier aims to help identify venomous and non-venomous snakes, using techniques of computer vision and machine learning. Their app uses a model of convolutional neural network trained with a dataset of 1,200 images of 12 species of Brazilian snakes. They have also used transfer learning. Different from us, they have used python, which makes it harder to create web application, our main selling point. They have used MobileNetV2, which is essentially what we have used.

1.2.3 Apps for biology

With recent technological advancements, several mechanisms have been employed to link technology with biodiversity. Many applications and programs have

utilized connections with local communities to bridge gaps in scientific knowledge related to biology, such as: recognition of sounds emitted by animals (e. g. FrogID, BirdNET, Merlin Bird ID; Rowley et al. (2023), Kahl et al. (2021)), birds identification (e. g. Merlin Bird ID), plant identification (e. g. PlantNet; Garcin et al. (2021)). In addition, some resources are produced directly for the scientific community, although there are instructions accessible to non-specialists (e. g. online multi-access interactive identification key; Pezzuti et al. (2021)).

Rowley et al. (2023) built a system for frog identification, which includes noise they make to support the identification. Merlin is an app for identifying birds, similar to ours. We have gathered more information on Merlin [here](#) since it is an interesting app. Basically, they are using the same approaches as we are: deep learning + transfer learning. INaturalist has on their website the possibility to try to identify the species on the photo: for Brazilian snakes, it does not seem to do a good job.

1.3 Article content

On this article, we do a test of concept, and our results are promising for a big model, a model that could from images, identify a snake. chatGPT has been incorporated into tools (Brockman, 2023) as so it can be used as an optimized way for teaching and learning, and we envision a possible integration. Given an image identified, we could use chatGPT as an addition aiding tool for teaching, gathering information, and making sure our model got it right.

We are going to consider three models. Since we are presenting a test of concept, those three models aim at showing different perspectives from the app. We believe that this is generic enough to be used in any similar case: one wants to separate images into classes (e.g., genus, species and subspecies). One just needs to change the input images and associate the respective class with each annotated image.

The first model classifies just false coral snakes, and it was the best results, achieving almost 100% of accuracy, with zero misclassification; and it is easy to converge, it does not require many attempts to get a good classification. However, when trying to classify with true coral snakes, the model starts to have problems (model 2) with very similar snakes, but still having accuracy higher than 50% (a measure of randomness). For the last model, model 3, it was added different species to see how the model handles diversity, for a possible big model (see Section 3); the true coral snakes were removed for avoiding the issue we already saw on model 2.

Finally, we present on the final section of the discussion section how we see a possible big model based on those findings. We are strongly convinced that a model for snake classification is not just possible, but also useful for the general public, and maybe for experts depending on the final accuracy. Apart from possible setbacks we may face on the way and we are optimistic we can solve them gradually, we strongly believe that building this app is imperative, and could contribute to the interception between biology and computer science (bioinformatics).

2 Methods

On this section, we present details about the training, and methodologies we have used. We have also used the opportunity to explain some basic concepts when needed.

2.1 Dataset

"Garbage in, Garbage out"

We have used images from Google Images: given the species/subspecies, we looked for images on Google Image. We had a couple of well-curated images for starting the model, but we have gradually added those extra images. Neural models, when well-trained, tend to ignore noise. Thus, even if some images may be mistagged on Google Images, it seems to be ignored since the training was a success. We hope to close a possible collaboration for retraining the model on a better dataset, curated by a biologist. However, what we got so far is promising.

We should also bring to attention that we already have a pretrained model for INaturalist available at TensorFlow Hub, one of the biggest datasets on species on the internet. This means that someone else has already trained a large model focused on biology, made available as feature model. [It can be accessed here for transfer learning.](#)

2.2 Training the model

We started with a small set of images, well-curated by a student of biology. Then, we considered the graphs, especially the matrix of accuracy per class, and added new images from Google Images when needed, and retrained the model. Even though we have changed the defaults parameters from Teachable Machine, the default parameters were enough. We have used extensively a feature in which they allow you to disable a class, which was useful to figure out which class was making the learning process unstable, and which classes were well-classified. In most of the cases, 0.001 as learning rate and 50 epochs were enough.

2.3 Neural networks, Deep learning and transfer learning

Neural networks are a subset of machine learning techniques (Fig. 1), focused on numeric algorithms, different from alternatives from artificial intelligence, those algorithms are not focused on reasoning; even though non-experts using those models may think they reason. Moreover, even though they are inspired by the brain, they do not replicate the dynamics of neurons with fidelity. Their only goal is learning from samples, and it does not matter how.

We are going to use the supervised variety: one presents a set of inputs, and expected outputs, and the algorithm should learn, without human interference, except at annotating the dataset.

On this type of algorithms, they should learn on their own, we do not interfere with the learning process.

Deep learning is a neural network with too many hidden layers; that is what 'deep' stands for, in opposition to




Figure 1: Neural networks are ruler finders. Source: Pires (2022a)

shallow, which now is used to designate classical neural models. They are designed to execute certain tasks, now called narrow neural networks, in opposition to artificial general intelligence.

Those models revolutionized the machine learning community, and are everywhere: from Google translator, chatGPT, DALL-E, to image classification. The drawback is that they are heavy to train, and require too many images (millions). The solution for making this technique widely accessible was transfer learning. In simple terms, you share your learning, and people can adapt to their small cases.

2.4 Training and validation

It is a common practice in the machine learning community to validate the models splitting the dataset into learning and validation: the ratio will appear on the figure for accuracy per class (e.g., Fig. 10). In all the scenarios we tested, the training curve converged.

When we say about 'instability' of the model, we are mentioning the validation curve. The validation curve is used to avoid memorization, overfitting. A model that memorizes is useless since it will have to predict outside their training dataset. Since our images are mainly from Google Image, when testing the model, we tried to make sure the testing was done with images not used neither on the training dataset nor on the validation dataset; if you input an image, say by accident, already present on the training/testing dataset, the model will give almost 100% of certainty on the prediction. The testing is just a confidence building strategy, it does not change the model.

A couple of samples is [here](#), for the general model.

Bear in mind that we are paying for a server on Heroku, those links may change in time. The reader is invited to get in touch by e-mail for the latest links and resources.

2.5 Teachable Machine

Teachable Machine is a Google AI Experiment, as they like to call it. It is built on top of TensorFlow.js. What you do with this platform can be built on TensorFlow.js, except that with this platform, no code needed. It does all the training ritual, from uploading image, preprocessing to showing training details. We have explored extensively this platform herein. In fact, we have taken advantage of that and we have actually done tests directly with

TensorFlow.js in order to compare the results: the results of Teachable machine is competitive with plain TensorFlow.js.

One possible usage of TensorFlow.js directly is actually testing on different feature models, with different feature numbers as output. You may see this feature number as a granulation from the model. The more features you have on the base model, the more details it is likely to find; see that we cannot control which features the model will extract, dictate where to look on the image. In essence, the feature model reduces the image into a smaller set of details, the features; see that other authors, as brought to attention on the introduction, actually used PCA, but the result was still better with feature models. Keep in mind that except for some final touch, the base model is not trained on the small dataset, it comes already trained.

For using Teachable Machine, one just needs to upload the images, and ask to train; it takes seconds, therefore, you can experiment with different parameters, or even retrain the model if it does not converge. The model is using local search, thus, it may get trapped on local minima. There are not too much to setup. We have kept the default parameters.

The one change we made sometimes was the training rate. The thumb rule is: small training rate will make the training slow, careful; whereas big values will make it faster. Think link this: when you want the model to be careful, try out small learning rates; when you believe the model can learn fast, just increase it. In most of the cases, the default parameters were the best. You may need to experiment with different training rates for either replicating our results, or adapting to your case.

2.6 TensorFlow.js

TensorFlow.js is a JavaScript-based library for deep learning, based on the classical TensorFlow, written in Python; you can also do simple learning machine, some simple mathematical operations with tensors and so on.

There are several reasons for using TensorFlow.js instead of Python, an imperative reason is using just one language, from the app to machine learning. Another reason is that you do the calculations on the browser, no need to have high-performance servers. If your app starts to gain users, the calculation cost will not grow since each user is responsible for their calculation load. This last point is especially interesting for startups, since you can scale up without also increasing the cost. One reason for medical applications: your data never leaves the browser, it is ideal for sensitive data.

A nice point is that they claim it is possible to transform models in both directions: TensorFlow.js < - > TensorFlow (Rivera, 2020; Shanqing Cai et al., 2020; Laborde, 2021). It is possible to transform models in both directions. Even the manual transformation is possible since their notations are similar.

TensorFlow.js provides several ways to be used: pretrained models, CDN calls, NPM, and even downloading the models. Each way has their own advantages and disadvantages, which we shall not cover here.

2.7 Feature model

Feature model (Fig. 2) is a deep learning trained model, but just the hidden layers: it is unable to make sense of an image since its classification layer was removed. It is made available as so others can use this knowledge to train small models, like ours: we just place a classification layer, and train the model on our datasets (i.e., snakes).




Figure 2: Scheme on how a model based on transfer learning works.

The biggest motivation for making those models available is that not everyone has access to high-performance computers. For our case, using machine learning in the browser, we have a limitation on how much we can take from a person's computer. See that TensorFlow.js also runs in the backend, using Node.js, which requires GPUs; in fact, since JavaScript is also present at desktop, you can also use on desktop apps. Another motivation is that having millions of images, even say 100 per class may be unreal. Transfer learning was created to tackle those limitations on actually using machine learning on images.

2.8 External resources

We cannot make sure those resource will be available when you read this article, but here goes the current public resource for learning from this article. You may get in touch by e-mail for latest resources and links. See that we have a [course teaching the basics of the paper](#).

Find [here](#) on GitHub how to use an already-trained model exported from Teachable Machine using Angular (TypeScript).

As an extra feature, teachable machine hosts your model, and make it available as a link, that can be called externally. It is ideal since we separate the model training and updates from the app; we have provided our links on the upcoming lines. Once you update the training, the change will automatically be pushed to the model being used on the link. If you prefer pure JavaScript, here

goes a [gist](#) on GitHub. If you need more options, which includes Android, you should access the export option of each model, after training.

Below the links for each model reported here. See that we had to retrain the models, thus, you may obtain different results. For the raw models, please, get in touch as so we can find a way to exchange the files.

- i. Fake coral snakes (model 1): https://teachablemachine.withgoogle.com/models/W9_qIu14Y/
- ii. Fake vs. true coral snakes model (model 2): <https://teachablemachine.withgoogle.com/models/9vw2M7LJw/>
- iii. Model with variety (model 3): <https://teachablemachine.withgoogle.com/models/Sc8mKQsS0/>

Tip. if you click on the link, the model will open on the browser and you can input images. The models of the paper will not be updated, as so you can test the results herein. The same links work if you are programmer, for importing the model inside your code.

3 Results and Discussion

On this section, we shall focus on the results achieved by running our model with a set of images. We have selected different snake species. This selection was done without any predefined preference, biased by the studies of the authors. We have a preference for Brazilian snakes, the field of studies of LHDB.

We cannot see how this selection would affect the results, thus, it does not affect reproducibility (as long as you use the same images, it likely to be reproducible).

We have used curated images alongside Google Image samples. By 'same images', we mean for the same classes: machine learning is not tied up to specific inputs, when properly developed.

3.1 model 1: just fake coral snakes

We have selected "randomly" four snakes classified as fake corals; numbers range from 50 to 80 different false coral snakes existent worldwide. This variation has to do with the observation that those classifications are not scientific, any snake that looks like a coral and it is not can be classified as false. Our goal is showing a scenario where two species or more are strongly alike, and how the model would behavior.

Namely: 1) *Apostolepis assimilis* (8 images); 2) *Erythrolamprus aesculapii* (18 images); 3) *Oxyrhopus rhombifer* (24 images); 4) *Lampropeltis triangulum triangulum* (LTT, 22 images).

We have used chatGPT to get a list as complete as possible from which we have selected the species used, and selected the correspondent images from Google Image; the list is [here](#). We have taken the care to avoid species with too many variations between subspecies: that is why we have *Lampropeltis triangulum triangulum* (subspecies) instead of just *Lampropeltis triangulum* (species).

Initial simulations with a possible big model showed that variations between subspecies can make the model unstable, and misclassify, generally amongst similar

species; also, the validation curve does not converge, and can even "blow up" when we place subspecies in the same class for the model to learn. Those classifications are constantly reviewed by the biology community, and subspecies may become species; we are using what is officially accepted at the time of writing this article.

In some applications, people are using chatGPT to automatically call apps from text (i.e., it understands what the user wants and given a set of apps, it calls those apps to give a well-based response), the results have been promising (Wolfram, 2023; Brockman, 2023). We have tested this in other scenarios, and the results are promising. As one example, chatGPT was able to suggest a snake from a short-text based description, even though it may get it wrong sometimes, and you may need to cross-references with other sources, it is a good starting point.

INaturalist is also able to identify from images, but, for Brazilian snakes, it seems to be bad at; sometimes it gets it right, but giving several suggestions which can be confusing.

Even the MobileNet (generally used as base model for transfer learning) can make an identification, but without details regarding species (it tags it as king snake, see Fig. 3); king snakes and corals are two different types of snakes with distinct characteristics; king snake is considered a false coral snake. Keep in mind that when the base model gets it right, it is transferred to the final model (i.e., the fact that MobileNet can spot a king snake does not make our model meaningless, our goal is making those models more accurate on their prediction, especially for Brazilian species).



Figure 3: King Snake, a false coral

Another possibility is Google image, however, adding to possible mistagging from users or even from Google when making "its trick", it tends to associate image with already-existent image; in fact, we had difficulties to find species from Google image, using the image as input. The user will most likely not take pictures from the snakes close to their images: chances are high the image will not exist, and chances are it will not necessarily be close to any existent one on Google Image. As one example how Google Image can get it totally wrong. We have entered a *papa lesma* (*Dipsas indica*): it gave several suggestions, some totally wrong, but did not get it right.

Our goal with transfer learning those big models is actually making them more specific. Since we are using Teachable Machine, we do not know the base model, it

is not available on their site. One alternative, which we do not explore here, is actually making transfer learning using TensorFlow.js manually, see Laborde (2021), chapter 11, for a guiding on how it is possible to use your own base model of preference. Even though harder than teachable machine, it may require programming abilities in JavaScript, it is still much easier than classical approaches using python.

Still on the base model, we have actually done a test with INaturalist: there is a pretrained model available as feature model that was trained on INaturalist. So far, the only advantage we saw is that the final model converges faster with INaturalist. Surprisingly, using INaturalist as base model did not improve the predictions. We have actually created the model outside Teachable Machine for comparison using JavaScript. More investigations will be done in the future for confirming this finding.

Generally, we have started to train with a small number of images, and increased when needed. We have used both the confusion matrix and the accuracy per class (Fig. 10) to decide which class to increase the image samples. What is interesting: in some of the simulations, the model converged even for 3 images. This is possible due to transfer learning, without transfer learning, a model would never converge with just say 30 images, unthinkable for 3 images. It also happens because they are very distinct snakes: the closer two species, the more images you may need.

We are going to present some results. The goal is showing a potential of a platform, a concept validation. We want to show that it is possible to build a model for snake classification with current free technologies, and possibly support on learning and even helping people to avoid accident with snakes.

Innovating with biomathematics is all about making our models easily available (Pires, 2022b). Usually, the scientific publications are the final goal on a research, but we believe in a culture where applications/innovations, should be the final goal. This is the startup culture.

Artificial intelligence, as APIs and libraries, is making this a reality, with models each time more powerful, and easier to use in generic applications.

Below the confusion matrix for model 1 (Fig. 4). This matrix shows that all the species were gotten right: no misclassification was done; when a misclassification is done, you should see values out of the main diagonal. Keep in mind it does not guarantee the model will not make mistakes. If the dataset is biased, say to a subspecies claiming to be a species, the model will behaviour like it learnt, but it will fail when presented with subspecies not represented on the training process.

Biased dataset is a big issue on machine learning: those models cannot learn what have not been shown. If they have never seen a species, it will try to guess amongst what it has.

As an example (Fig. 5), the model 2 is misclassifying, which has two true coral snakes, with 60% the *micrurus lemniscatus* which does not exist on our dataset at the moment as *micrurus frontalis* (true coral) and 36% as *Erythrolamprus aesculapii* (the winner on misclassifications on the current "big model" since it is quite similar to the true corals). See that we had to retrain




Figure 4: confusion matrix for model 1

the model to run this test. Always keep in mind that you may obtain results slightly different from ours since those models are initiated with random weights.

Figure 5: *Micrurus lemniscatus*, true coral not present at the current training. Source: Wikimedia Commons

To make sure we are not memorizing the dataset, you can check out Fig. 6, which shows that both learning and validation curves converged to lower values; and best of all, the validation curve is slightly pointing downward, which means we could improve the model by increasing the learning time. See that teachable machine allows to stop the learning process, at any time: you may add an big learning time, and stop when you see it is becoming unstable, or converged.

Another useful graph is Fig. 7, which presents how much the model got it right: training arrived to 100%, and validation to 99%; it is called accuracy, you divide the correct classifications by the total attempts. All those curves help to avoid common traps when teaching those models (e.g., overfitting).

“With four parameters I can fit an elephant, and with five I can make him wiggle his trunk” Johnny von Neumann

As a last way to make sure the model is working, we can




Figure 6: learning curves for model 1. Legend: the blue curve is training, whereas the orange one is testing/validation.




Figure 7: Accuracy for model 1.

actually test it (Fig. 8); we have actually applied on a more complex image, which also worked, a person holding it on his/her hand, but due to possible copyright issues, we have decided to show this simple case as illustration. The model gave 80% of chance of being an *oxyrhopus rhombifer*, with lower probabilities for the remaining classes; high values for remaining classes even when the model gets it right show possible confusion, chances of misclassification. It is a good result for classification. see also Fig. 9. Remember that the model will give probabilities for each class present at the training, and we choose the highest one. We are playing with *clustering the probabilities* by genus; and also to use extra information such as geographical.

Still on Fig. 8, different from Merlin, which is focused on birds, it is not necessary to zoom in: it is done automatically by the algorithm. We have though about it, however, if we decide to use zoom in, we are planning to use machine learning to automatically zoom in. Unless we have no choice, it is interesting to avoid human's



Figure 8: Oxyrhopus Rhombifer. Source: Wikimedia Commons



Figure 9: Realistic photo for *erythrolamprus aesculapii* (71% given by the model, and additionally with low probability to other similar classes. See that this image was chosen for being hard to classify due too much background and obstructions, being very realistic).

interference.

Last but not least, Fig. 10 shows the classification per class. Since the confusion matrix was diagonal, we already expected it.

3.2 model 2: false coral snakes vs. true ones

We have more than 50 true coral snakes, and even more when it comes to the false ones; keep in mind that it is not a scientific classification. The numbers are not a common consensus; for the false ones, the subspecies may increase

Accuracy per class

CLASS	ACCURACY	# SAMPLES
Apostolepis assimi...	1.00	2
Erythrolamprus aes...	1.00	3
Oxyrhopus rhombife...	1.00	4
LTT	1.00	4

Figure 10: accuracy matrix per class for model 1

even more the final count.

On this model, we are using the following true coral snakes: 1) *Micrurus corallinus* (24 images); 2) *Micrurus frontalis* (25 images); 3) *Micrurus alleni* (23 images); 4) *Micrurus albicinctus* (14 images). For the false ones: 1) *Apostolepis assimilis* (8 images); 2) *Erythrolamprus aesculapii* (18 images); 3) *Oxyrhopus rhombifer* (24 images).

This is a tough classification. Humans generally cannot make the difference, unless, they are experts, with specialized training. This paper was born because JGP could spot a false coral, LHDB can due his technical training in biology.

This model can be useful for helping people to separate the fake coral snakes from the true ones. The fake ones are generally not dangerous (i.e., venomous snakes). The true coral snakes are well-known for its potent venom, it can be fatal to humans.

Error type I and error type II are two different types of errors in statistics: they are called false positive and false negative. We are calling on this model error type I as classifying a non-dangerous snake as dangerous; and, similarly, error type II as classifying a dangerous as non-dangerous. See that even though one can use our models to separate venomous snakes from non-venomous ones, different from Santos et al. (2021), this is not our primary goal.

We assume we all agree that error type II should be avoided, if we assume one key application is actually helping people avoiding accident with those snakes. On the current version of the algorithm we cannot make sure we avoid one type of error, but we can see the final training, and retrain until we get an expected result. For this case, it may take several attempts, which takes seconds to train.

A possible enhancement of the model is considering automatization of this best model finding. We see promising results actually trying an evolutionary computing. We are able to find even already available repositories using TensorFlow.js with genetic algorithms, something we need to study and run simulations.

Different from model 1, this model did not converge easily. it took several attempts to get the results showed here. The model has problems to classify similar snakes. One example was *Micrurus corallinus* vs. *Erythrolamprus aesculapii* (Fig. 11); *erythrolamprus aesculapii* is the champion of fooling our model, and humans too (Fig. 11). We saw similar problems with tests with a model not



Micrurus corallinus



Erythrolamprus aesculapii

Wikimedia Commons

Figure 11: Micrurus corallinus vs. Erythrolamprus aesculapii

reported here for just *Bothrops* genus: similar snakes tend to induce misclassification between themselves (Fig. 14).




Figure 12: Confusion matrix for model 2

On Fig. 11, just take a look at the *Erythrolamprus aesculapii* on horizontal line, and look upward (basically, all the misclassifications come from *Erythrolamprus aesculapii* being misplaced as true corals). Similarly, do the same on the vertical, and also the misclassifications come from *Erythrolamprus aesculapii* being misplaced as true corals. The errors type II, to be avoided, are on the lower part of the diagonal. For reading this matrix, on y-axis (vertical), we have the true class; whereas, on x-axis (horizontal), we have the model predictions (Fig. 12).




Figure 13: Accuracy for model 2

See that if you try to replicate those results (Fig. 13), you may find something different, likely similar. Neural models are initiated randomly, thus, when you start training, chances are you will start on a different point on the neural network space that we started. Just try a couple of times, and the model should converge as we report (Fig. 15). Fig. 16 shows accuracy per class. See that most of the misclassifications come from *erythrolamprus aesculapii* being misplaced as either *micrurus frontalis* or *micrurus micrurus corallinus* (both dangerous). We are not sure adding more image will help; even though we may try in the future.

On the machine learning community, we know those models are generally bad at what we humans do easily; and tend to be better at what we do poorly. It seems biologists make the difference on those cases the model



Figure 14: Bothrops alternatus misclassified as Bothrops neuwiedi. Source: courtesy of Luiz Braga




Figure 15: Training curves for model 2

Accuracy per class

CLASS	ACCURACY	# SAMPLES
Apostolepis assimi...	1.00	4
Erythrolamprus aes...	0.60	5
Oxyrhopus rhombife...	0.75	4
Micrurus corallinu...	0.50	4
Micrurus frontalis	0.50	4
Micrurus alleni	1.00	4
Micrurus albicinct...	1.00	3

Figure 16: Accuracy for model 2

is struggling with by taking a look at the tail, head and belly. chatGPT is suffering from "common sense" (easy answers for humans, and it makes a mess when trying to answer)(Choi, 2023). Thus, what we human created "a common sense", the model cannot see. The model used as base, for transfer learning, is basically frozen on our small dataset. Therefore, we cannot be sure it learnt to actually "look at tails" or head. How those models learn is something we cannot explain, chatGPT has been target of several nice discussions (Wolfram, 2023). Those discussions are not just helping us to popularize AI, but also to question basic assumptions we have been doing. chatGPT opened the gate for talking about artificial

general intelligence, smarter models.

3.3 model 3: experimenting with diversity

In the world, 4,038 different species of snakes are listed Uetz et al. (2022). 435 of these species are registered for Brazil, being the 3rd country in the ranking of the highest number of snake species. For building such a model, it would take time, patience, and lot of collaborations. Our simulations shows that when the snakes are similar, not necessarily from the same genus, they tend to be misclassified amongst themselves.

We are not even sure it is possible to build such model, at least, not with 100% of accuracy: mistakes are about to be made on such a generic model. The best model we know is able to identify something superior to 1.000 different images. MobileNet, widely used in transfer learning, can identify 1.000 different objects; even though, they can make mistakes, or can "get in doubt" ("confused"). The Inception model also can identify this number of objects Laborde (2021). They differ on precision and speed of response. See that those numbers may change since those models are constantly enhanced. A solution would be building small models, with species separated by rules (e.g., fake coral vs. true ones). The selection of the model to be used can be done by an algorithm, by the user, or even by chatGPT; something we must experiment on.

A good result we can see on the literature, which we are not sure it is easily replicable is Kalinathan et al. (2021): they created a model with more than 700 different snakes. Thus, it may be possible to build such a model, they used "old" approaches, we have better methods now.

For the fake ones: 1) *Apostolepis assimilis* (8 images); 2) *Erythrolamprus aesculapii* (18 images); 3) *Oxyrhopus rhombifer* (24 images). And for the diversity, we have used: 1) *Crotalus durissus* (12 images); 2) *Bothrops alternatus* (24 images); 3) *Bothrops neuwiedi* (22 images); 4) *Bothrops jararaca* (23 images). Since we already know the difficulties between false and true coral snakes, we left the true coral snakes out.

Accuracy per class

CLASS	ACCURACY	# SAMPLES
Oxyrhopus rhombife...	1.00	2
Erythrolamprus aes...	1.00	3
Apostolepis assimi...	1.00	4
Bothrops jararaca	0.50	4
Bothrops neuwiedi	0.50	4
Bothrops alternatu...	1.00	4
Crotalus durissus	1.00	2

Figure 17: Accuracy per class for model 3. Legend: 1 – 100%, no mistakes; 0 – 0%, got it all wrong.

As we can see from Fig. 17, the accuracy per class ranged from 50% to 100%. Another interesting fact: a small model tends to be good also on the big model (see that no fake coral snake was misclassified), which something we need to investigate. It would be good news if that observation holds for all the classes.

The mistakes accounts for misclassifications on the *Bothrops* genus (Fig. 19). We have a guess that part of this misclassification may come from the fact that they have subspecies, which we have verified, with wide variations




Figure 18: Accuracy for model 3

in patterns: we may need to train for subspecies.




Figure 19: Confusion matrix for model 3. Legend: the numbers mean how many classifications are right, when on the main diagonal, or misclassification, when outside the main diagonal. The perfect matrix is diagonal.

The good news is: the model misclassify amongst the same genus (Fig. 21). The solution we found for now: the model tend to give high probabilities to similar snakes. We just sum up the probabilities for genus, and add this as the predicted genus. We call this *probability clustering*.

It may help for a potential second level where we give the first prediction for a more refined model, something we want to experiment on for improving the predictions. This strategy could also increase the number of species/subspecies we could spot.

As you could see, the models can misclassify (and a potential big model, Fig. 22), which generally happens when the snakes are similar. *Erythrolamprus aesculapii*




Figure 20: Training curves for model 3

was the winner, able to fool our model, passing as true coral snakes (Fig. 12).

As we highlighted before, the dangerous is when the venomous snake is classified as false coral snake, what we have called error type II. The images show that they are very similar in patterns (Fig. 11). A solution that came up during our discussions is using geographic information. The collaborator on this research is actually working as part of his research on a probability distribution of snakes based on geographic information. However, we are opened for possible collaborations. We have done some tests with INaturalist API, that provides this information once we provide a location. It could be used to support on predicting the snake: in addition to image identification, also considering whether the snake was ever seen on the user's location.

As alternative, we may consider public APIs that provide information and more. INaturalist has an API, which provides information based on API calls. Providing extra information, we could actually give our model a better chance to get it right. Another alternative is actually asking the user extra information, either before or after the prediction; we are avoiding this type of strategy as so the system will be as autonomous as possible. We should mention that chatGPT (Fig. 22) showed promising behaviors during this app development as an assistant, which could be used also to enhance the predictions (e.g., chatGPT is able to guess, sometimes getting close, the snake based on a short textual description).

3.4 Conclusion

We have presented a state-of-the-art machine learning system for classifying snakes, best of all, no coding during the model development. It was possible thanks to a public platform by Google called Teachable Machine. This approach can be adapted to any classification problem, as long as the difference between the species is on the images. Our hope is that we do not just call attention to this option for non-experts in machine learning, but also call for collaborators, especially, with well-curated datasets




Figure 21: Misclassification of the species, but classified properly the genus. This case was taken from our actually deployed app

on snakes species.

It is hard to make a fair comparison with the other authors cited herein. The reason is that those models evolve very fast, therefore, a better result here may be due to a better routine for image identification. One strong point of our model is that it was implemented basically codeless, using an online tool. It makes it easier to replicate the results by non-experts in programming (e.g., biologists), or even, for passing the work forward. Also, the training is decoupled from the app, therefore, it is possible work on them separately. It will certainly make it easier to share the work load on possible collaborations.

See that we have an online course teaching the basics of this paper for the general public. The best result we found on the literature was able to classify +700 different species, at least theoretically, we can arrive to 1.000, and with some tricks, even more. One challenge is that a species may have several subspecies, on some cases, those subspecies are not necessarily alike in the image. Some species may have too much variations, which make it harder to use species as an way to enhance classification (i.e. clustering the probabilities).

One insight we have mentioned on the paper is using geographic information to support on the classification. Thus, if a snake is classified as probable, we can see if the snake was ever seen on the geographic location it was found by the user. This information is part of the work of LHDB. For creating an initial prototype, we can use the iNaturalist API. It allows a call with location, and it gives back several pieces of information, amongst this




Figure 22: Vision of a possible execution plan for SnakeFace based on the presented results

information, we have the frequency the species was seen on a location. This information can be crossed with the model prediction to improve the prediction. Another application of INaturalist, a huge dataset on species, is using their feature model, already available on TensorFlow Hub. We did some initial testings and comparisons: we could not find so far advantages compared to what we did so far. We hope to make more testing.

In one of the papers mentioned, they have used PCA and compared with an approach similar to ours, using transfer learning: transfer learning won. It means that most likely we are on the right track.

One possible drawback is that we are using JavaScript, since python is the dominant language in machine learning. It is not necessarily a problem since we can convert the models between the languages, thus, if you train in python, you can just convert the model to JavaScript. Therefore, if a new model is released, and it is in python, we can still make the testing. It is possible that new models will be better at predicting, and will have more objects (so far, the best is about 1.000 different objects). Each species/subspecies is an object. See that we have released [our datasets on Kaggle](#), a python machine learning community, so far, no one has solved the problem as we did using python.

Our biggest challenge now is actually adding new species, and keep training. It may take time, and at the moment we do not have funding. As we train, we need to study the curves, and make sure the model is optimized for a certain goal. We mentioned about error I and error II:

we should avoid misclassifying poisonous snakes.

We never know for sure how our findings will be applied. One futuristic application of our findings is when integrated with drones (i.e., their webcam). As far as we know, drones cannot move around areas full of obstacles (e.g., wide forests), something that may change fast; even the fact that we are making this proposal herein could create the motivation for such researches, that could focus on that. We are aware of researches that try to use drones to spot targets on physical areas. Once a drone is able to move in a wild forest, one could sample all the snake species in an area, even look for a specific species. No need of human's interference, avoiding possible accidents, or even human's interfering with the wide life. Teachable Machine allows to choose between photos and webcam, also, it is already a standard practice between TensorFlow.js practitioners, the library behind Teachable Machine, to use webcam for identifying objects. One just need to replace the computer webcam with the drone's one.

4 Author contributions

JGP created the model, ran the simulations, designed the codes, and wrote the manuscript. LHDB helped with the snake classifications, initial images, writing the manuscript, and supported on designing the software to best fit biologists' requirements for being friendly.

References

- Ahmed, K., Gad, M. A. and Aboutabl, A. E. (2023). Snake species classification using deep learning techniques, *Multimedia Tools and Applications* . <https://doi.org/10.1007/s11042-023-16773-0>.
- Almeida, P. C. R., Feitosa, D. T., Terribile, L. C., Freitas-Oliveira, R. and Prudente, A. L. C. (2022). Co-occurrence patterns between false coral snake *Atractus latifrons* (günther, 1868) (serpentes: Dipsadidae) and venomous coral snakes from the amazon, *Anais da Academia Brasileira de Ciências* **94**(Suppl. 3). <https://doi.org/10.1590/0001-376520220210991>.
- Altman, S. (2023). *OpenAI CEO on GPT-4, ChatGPT, and the Future of AI*. Available at https://www.youtube.com/watch?v=L_Guz73e6fw&t=7260s.
- Brockman, G. (2023). *The Inside Story of ChatGPT's Astonishing Potential*, TED. Available at https://www.youtube.com/watch?v=C_78DM8fG6E&t=626s.
- Campbell, J. A. and Lamar, W. W. (1989). *The Venomous Reptiles of Latin America*, Comstock Pub. Associates.
- Choi, Y. (2023). *Why AI Is Incredibly Smart and Shockingly Stupid*, TED. Available at <https://www.youtube.com/watch?v=SvBR00GT5VI>.
- Davis Rabosky, A. R., Cox, C. L., Rabosky, D. L., Title, P. O., Holmes, I. A., Feldman, A. and McGuire, J. A. (2016). Coral snakes predict the evolution of mimicry across new world snakes, *Nature Communications* . <https://doi.org/10.1038/ncomms11484>.
- Durso, A. M., Moorthy, G. K., Mohanty, S. P., Bolon, I., Salathé, M. and Ruiz de Castañeda, R. (2021). Supervised learning computer vision benchmark for snake species identification from photographs: Implications for herpetology and global health, *Frontiers in Artificial Intelligence* **4**. <https://doi.org/10.3389/frai.2021.582110>.
- FRONTLINE (2020). *In the Age of AI*, FRONTLINE. Available at https://www.youtube.com/watch?v=5dZ_lvDgevk.
- Garcin, C., Joly, A., Bonnet, P., Affouard, A., Lombardo, J., Chouet, M., Servajean, M., Lorieul, T. and Salmon, J. (2021). Pl@ntnet-300k: a plant image dataset with high label ambiguity and a long-tailed distribution, *NeurIPS Datasets and Benchmarks 2021*. Available at https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/7e7757b1e12abcb736ab9a754ffb617a-Paper-round2.pdf.
- Greene, H. and McDiarmid, R. (1981). Coral snake mimicry: does it occur?, *Science* **213**: 1207–1212.
- Hossie, T. J. and Sherratt, T. N. (2013). Defensive posture and eyespots deter avian predators from attacking caterpillar models, *Animal Behaviour* . <https://doi.org/10.1016/j.anbehav.2013.05.029>.
- Kahl, S., Wood, C. M., Eibl, M. and Klinck, H. (2021). Birdnet: A deep learning solution for avian diversity monitoring, *Ecological Informatics* **61**: 101236. <https://doi.org/10.1016/j.ecoinf.2021.101236>.
- Kalinathan, L., Balasundaram1, P., Ganesh, P., Bathala, S. S. and Mukesh, R. K. (2021). Automatic snake classification using deep learning algorithm, *Conference and Labs of the Evaluation Forum*, . Available at <https://ceur-ws.org/Vol-2936/paper-135.pdf>.
- Laborde, G. (2021). *Learning Tensorflow.js: Powerful Machine Learning in JavaScript*, O'Reilly Media. Available at <https://www.amazon.com.br/Learning-Tensorflow-Js-Powerful-Machine-JavaScript/dp/1492090794>.
- Pezzuti, T. L., Leite, F. S. F., Rossa-Feres, D. d. C. and Garcia, P. C. A. (2021). The tadpoles of the iron quadrangle, southeastern brazil: A baseline for larval knowledge and anuran conservation in a diverse and threatened region, *South American Journal of Herpetology* **22**(sp1): 1–107. <https://doi.org/10.2994/SAJH-D-20-00042.1>.
- Pires, J. G. (2018). *Relatório Final de pós-doutorado Programa Nacional de Pós-doutorado PNPd/CAPES*. Available at https://www.researchgate.net/publication/329815289_Relatorio_Final_de_pos-doutorado_Programa_Nacional_de_Pos-doutorado_PNPdCAPES.
- Pires, J. G. (2022a). *Computational Thinking: How computers think, decide and learn, when human limits start and computers champ. Vol.1*, self published, Kindle Version. Available at <https://www.amazon.com/Computational-Thinking-computers-decide-limits-ebook/dp/B0B5XZCVCL>.
- Pires, J. G. (2022b). Innovating with biomathematics: the challenge of building user-friendly interfaces for computational biology, *Academia Letters* . <https://doi.org/10.20935/AL5792>.
- Pires, J. G., da Silva, G. F., Weyssow, T., Conforte, A. J., Pagnoncelli, D., da Silva, F. A. B. and Carels, N. (2021). Galaxy and mean stack to create a user-friendly workflow for the rational optimization of cancer chemotherapy, *Frontiers in Genetics* **12**. <https://doi.org/10.3389/fgene.2021.624259>.
- Poulton, E. B. (1890). *The colours of animals: their meaning and use, especially considered in the case of insects*, D. Appleton and Company.
- Rajabizadeh, M. and Rezghi, M. (2021). A comparative study on image-based snake identification using machine learning, *Scientific Reports* **11**(1). <https://doi.org/10.1038/s41598-021-96031-1>.
- Rivera, J. D. D. S. (2020). *Practical TensorFlow.js: Deep Learning in Web App Development*, Apress. Available at <https://www.amazon.com/Practical-TensorFlow-js-Deep-Learning-Development/dp/1484262727>.
- Rowley, J., Callaghan, C., Cutajar, T., Portway, C., P. K., Mahony, S., Trembath, D., Flemons, P. and Woods, A. (2023). Frogid: Citizen scientists provide validated biodiversity data on frogs of australia, *Herpetological Conservation and Biology* **14**(1): 155–170. Available at https://www.herpconbio.org/Volume_14/Issue_1/Rowley_et_al_2019.pdf.
- Roze, J. A. (1996). *Coral snakes of the Americas: biology, identification, and venoms*, Krieger.

- Ruxton, G. D., Sherratt, T. N. and Speed, M. P. (2004). *Avoiding Attack: The Evolutionary Ecology of Crypsis, Warning Signals and Mimicry*, Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198528609.001.0001>.
- Santos, F. G. D., Machiaveli, E. A. and Marciano, J. M. V. (2021). Snake classifier: Aplicativo mobile para classificação de serpentes peçonhentas. Available at https://www.academia.edu/95702218/Snake_Classifier_Aplicativo_mobile_para_classifica%C3%A7%C3%A3o_de_serpentes_pe%C3%A7onhentas.
- Savage, J. M. and Slowinski, J. B. (1992). The colouration of the venomous coral snakes (family elapidae) and their mimics (families aniliidae and colubridae), *Biological Journal of the Linnean Society* 45: 235–254. <https://doi.org/10.1111/j.1095-8312.1992.tb00642.x>.
- Secretaria de Saúde (2022). *Acidentes por Serpentes*, Paraná, governo do estado, secretaria de saúde. Available at <https://www.saude.pr.gov.br/Pagina/Acidentes-por-Serpentes>.
- Shanqing Cai, S. B., Nielsen, E. D. and Chollet, F. (2020). *Deep Learning with JavaScript: Neural networks in TensorFlow.js*, Manning. Available at <https://www.amazon.com/Deep-Learning-JavaScript-networks-TensorFlow-js/dp/1617296171>.
- Stack Overflow (2023). *2023 Developer Survey: Programming, scripting, and markup languages*, Stack Overflow. Available at <https://survey.stackoverflow.co/2023/#programming-scripting-and-markup-languages>.
- Uetz, P., Freed, P., Aguilar, R. and Hošek, J. (2022). The reptile database. Available at <http://www.reptile-database.org>.
- Wickler, W. (1968). *Mimicry in Plants and Animals*, McGraw-Hill.
- Wolfram, S. (2023). *What Is ChatGPT Doing ... and Why Does It Work?*, Wolfram Research, <https://www.amazon.com/What-ChatGPT-Doing-Does-Work/dp/1579550819>. Available at <https://www.amazon.com/What-ChatGPT-Doing-Does-Work/dp/1579550819>.
- Yang, Z. and Sinnott, R. O. (2021). Snake detection and classification using deep learning, *Proceedings of the 54th Hawaii International Conference on System Sciences*, Proceedings of the 54th Hawaii International Conference on System Sciences, pp. 1212–1222. Available at <https://scholarspace.manoa.hawaii.edu/server/api/core/bitstreams/f9f2de4d-ddaa-4984-9fc3-9bbca4e27247/content>.