Seeing spots and counting chirps

Ian Durbach tells us how machine learning can be used in ecological surveys

Counting animals is an important part of many kinds of ecological problems. Ecologists are often interested in what makes one area better for some animals than others, a question that can only be answered after counting how many animals live in each area. And to assess whether a conservation project in a national park is working, managers would need to know if the number of animals in the park is going up or down.

The two main ways animals are counted is by seeing or hearing them. Both of these are difficult jobs, requiring a lot of training and many hours of fieldwork. Up until a few years ago, the most common way of doing this surveying work would be to have highly trained specialists go out and try to see or hear animals. Naturally, there were only so many of these specialists, which meant that the number of studies that could be done was fairly small.

Recent advances in digital camera and audio recording technology offer a different way of collecting data. Instead of people doing the counts, it is now possible to set up cameras or microphones, and let these do the work. But normally someone still needs to sort through the photographs, or listen to the recordings, to identify which ones contain the animals that are being studied. For example, the cameras that are used are usually motion-sensitive, and will take a photograph whenever any kind of motion is detected. This can include the species of interest, but also other species, or even motion that is of no interest at all, such as leaves moving in the wind. Audio recording equipment is usually set to switch on for certain hours of the day, and will record all sounds in the nearby environment, whether the animal is there or not.

Sifting through all this data can be very time-consuming, and these days data is being captured faster than it can be manually processed. Machine learning plays an important role in ecology, by developing models that can automatically recognise animals in photographs and audio recordings. These are models that recognise patterns – like a leopard's spots, a giraffe's long neck, or a frog's croak or chirp – that distinguish what we are interested in (the animal) from the rest of the background.

There are a number of different kinds of identification problem that can be assisted by machine learning. One is individual animal identification. Some animals have unique markings that allow individual members of the same species to be told apart. The most famous example is the leopard, but many species of frogs and toads also have unique markings. This is important when counting animals from photographs because there might be

Alex Rebelo

Distinctive patterns on western leopard toads allow individuals to be identified and counted.

The Cape Peninsula moss frog lives in moist habitats near streams and seeps. Adults don't grow much larger than 2 cm.

more than one photograph for some of the animals, so population sizes will be overestimated if this is not noted.

We work with one species of toad like this, the western leopard toad. The toads have striking patterns on their backs that are unique to each toad. We use a special model called a Siamese convolutional neural network to predict whether two photographs are from the same individual or two different individuals. The model is built by first feeding it pairs of images that come from the same individual, and then feeding it pairs of images from different individuals. From these examples, the model learns which features to pay attention to, and which can be ignored. For example, the model would learn that it is a toad's spots that make it different from any other toad, rather than – say – the number of legs it has.

Cape Peninsula moss frogs present a different kind of challenge. These frogs are extremely difficult to find. Despite being common on Cape Town's peninsula, very few people have seen one. They are only 2 cm long, live in thick vegetation around mountain seeps or streams, and stop calling as soon as approached. The only possible way of counting the frogs is by setting up microphones and recording how many individual chirps there are. To get from a count of chirps to a count of frogs, we need to divide by the average number of calls made by each frog – fortunately something that is reasonably well known.

Machine learning helps here by providing a model that is able to 'listen' to thousands of hours of recordings and very rapidly count how many chirps are present. It does this by taking one tiny slice of recording at a time – roughly 1/1000 of a second – and predicting whether there is a chirp in that segment or not. It learns, like the western leopard toad model, by example. We first need to feed the model lots of segments that do contain a chirp. Then we feed it segments that do not. Over time, the model learns to recognise what makes the two types of segment different.

These are just two examples where machine learning is allowing ecologists to focus on the more important scientific questions arising from their research, rather than spending many hours on mundane tasks like sifting through photographs or listening to recordings. As drone technology and high-resolution video and audio recorders become more affordable, ecological monitoring programmes are collecting more data than ever before. Machine learning is almost certain to play an increasingly important role in ecology in the years to come.

Dr Ian Durbach spends half the year in Scotland, where he is a research fellow at the Centre for Research into Ecological and Environmental Modelling at the University of St Andrews. For the other half, he is an associate professor in the University of Cape Town's Department of Statistical Sciences and an associate research fellow at AIMS. The research reported here was done by two UCT Master's students he supervised, Emmanuel Kabuga (leopard toads) and Jenicca Poongavanan (moss frogs).