

Object Detection Labelling Instructions (v 1.0)

Full run, Jan 2020


Authors: Emily Muller and Ricky Nathvani
r.nathvani@imperial.ac.uk

Goal

This document includes an overview and instructions on how to perform the image labelling and validation, for producing a dataset as part of the object detection efforts of the Pathways project. This is to assist and **ensure consistency** amongst labellers when performing the task to produce **high quality labels** for the data. Any questions about the instructions provided in this document should be emailed to r.nathvani@imperial.ac.uk

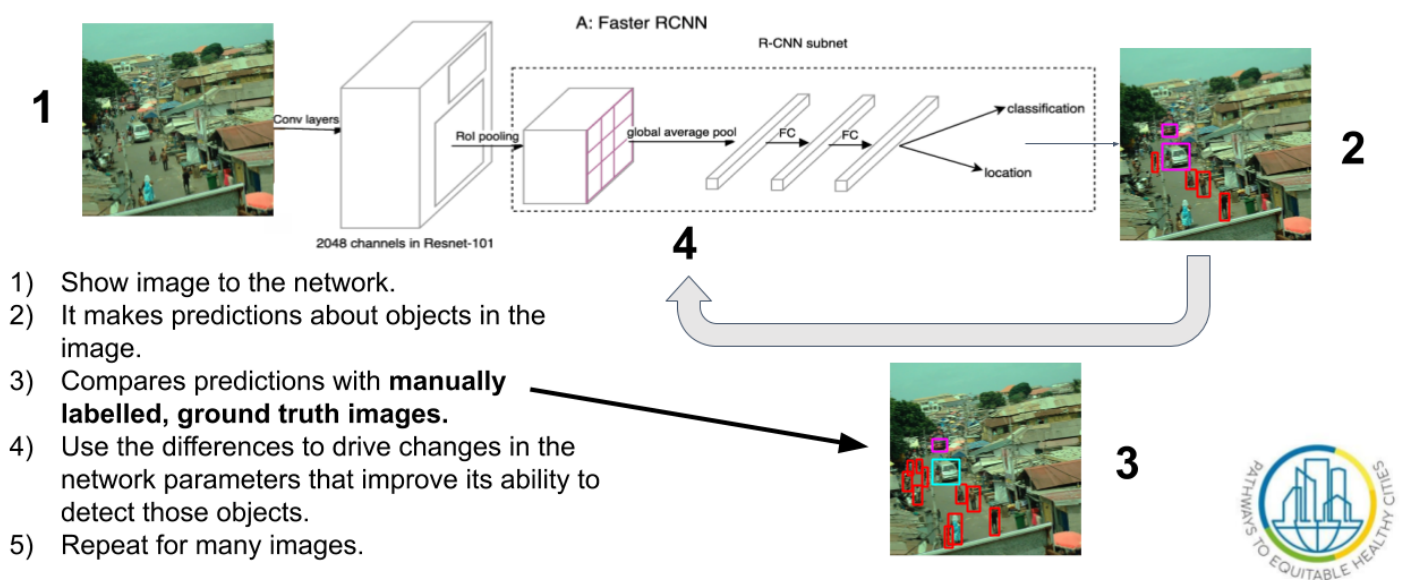
Object Detection: recap

- The goal of **Object Detection** is to locate the presence of objects (specifically “things” that can be identified in unique instances with well defined boundaries, such as footballs and coffee mugs **unlike** “stuff” such as clouds or grass) in images and surround each instance with a tight **bounding box**, as in the illustration below.



- This is achieved using a **machine learning** algorithm, which takes an image as an input and produces the categorised boxes as an output. Machine learning algorithms learn from **training data**, which are “correct” examples of the sorts of outputs we want to produce, that can successively improve the model to produce a similar output when presented with the same input, as shown in the diagram below:

Training an OD algorithm



- For an Object Detection algorithm, the training data consists of images where the boxes have been **manually labelled** by human labellers. As part of the imagery collection campaign in Accra, the Pathways project is looking to make a set of training data of exactly this sort by having a team of labellers put boxes around **every** object of interest in **a sample (not all!)** of images acquired from the Accra campaign.
- This will require each labeller to **label a certain number of images in line with the guidelines provided below**.

Object detection labelling

The core task is simply to label **every** instance of **a given object class** with a box in an image. After an extensive shortlisting performed in consultation with many Pathways members, the object classes we are labelling for the full data run are:

- Cars
- People
- Trotros
- Market stalls
- Trucks
- Cookstoves
- Motorcycles
- Street vendors
- Lorries
- Umbrellas
- Buses
- Trash
- Taxis
- Vans
- Pieces of Debris
- Loudspeakers
- Cooking bowls
- Food
- Animal
- Bicycles

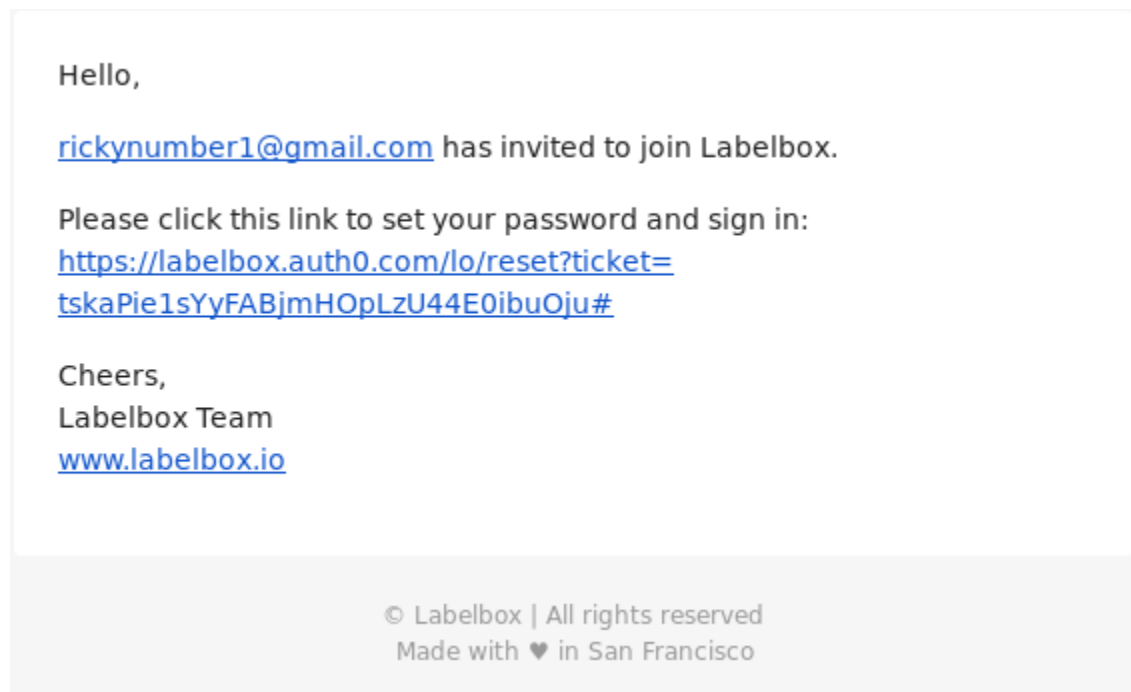
Therefore, each labeller should look to label **every instance** of **every object listed above** in their sample of images. Examples of all of these objects in the imagery and some examples of how to label them are provided in the following document, which **every labeller should read before doing any labelling to ensure labels are consistent between labellers**: [Object Detection Labeling Examples](#)

During the full data run, each labeller is expected to label about **20 images per working week** and about **150 images** overall. There is a way to keep track of this number as you work as explained below. Although 20 images may not seem like very many, and some images will be very sparse in objects making them easy to label, we emphasise that **quality** labels with good bounding boxes and sufficient **coverage** may take some time.

As explained above, the primary task of data labelling is drawing a box with your computer cursor around each visible object of interest in an image.

The software we have chosen to use, to enable labellers to perform this task, is a website called **Labelbox**.

You should have received an email (if not, please contact us at r.nathvani@imperial.ac.uk) with the subject line "Invitation to join Labelbox", with the following:



Follow the link provided to set up a LabelBox account. **Note:** If you set up your account with your gmail address but provide a separate password for your Labelbox account, choosing the option to sign in with your Google account in the login screen:



may not give you access to the project. Instead, you may have to log in with the company email and password prompts in the login page:

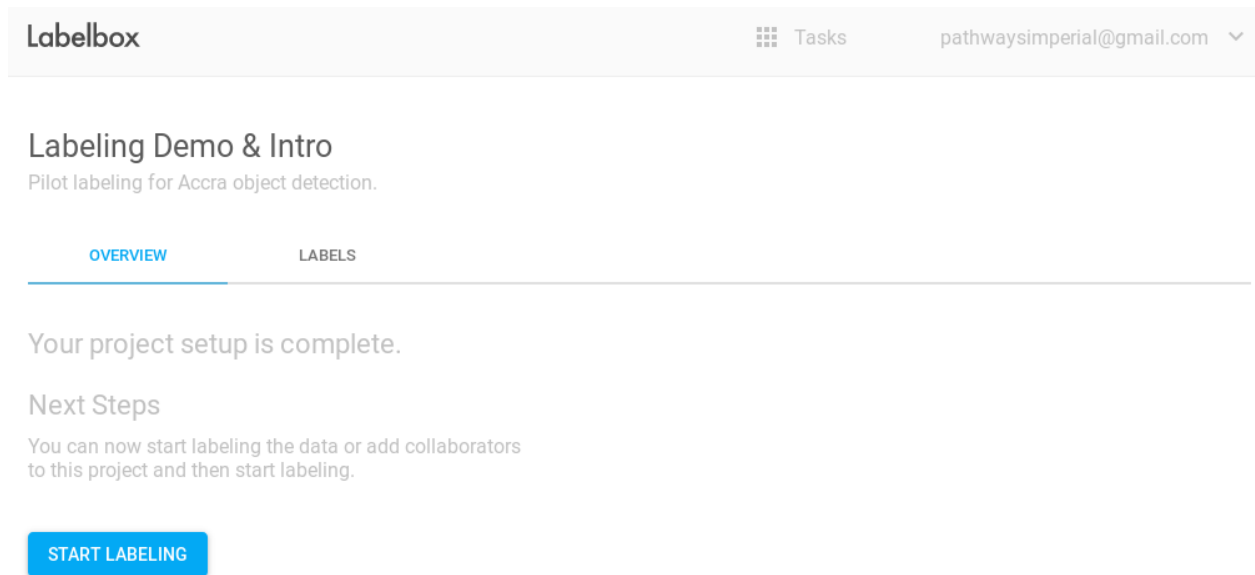
or

The image shows two input fields stacked vertically. The top field has an envelope icon on the left and the placeholder text "your company email". The bottom field has a lock icon on the left, the placeholder text "your password", and an eye icon on the right to toggle password visibility.


where you can supply the email and password combination you have registered with. Once you log in, you will land on the following page:

The screenshot shows the Labelbox dashboard. At the top left is the "Labelbox" logo. To its right are a grid icon and the word "Tasks". Further right is the user's email "pathwaysimperial@gmail.com" with a dropdown arrow. Below this is a blue "PROJECTS" header. Underneath is a search bar with a magnifying glass icon and the word "Search". At the bottom, there is a project card for "Labeling Demo & Intro" with a circular icon containing the letter "L". The card text includes "Pilot labeling for Accra object detection." and "0 labels Active a few seconds ago".

There may be some other projects visible, but the ones for the Pathways work are called **Accra Object Detection Labels (NOT Labeling Demo & Intro, shown in the examples)**. Just click on the name of the project to get access to it. Once there, you'll get the following page:




Click on the blue “start labelling” button to access the labelling interface (where the task is carried out). It will look something like this, though the image will probably be a different one:




On the left hand side of the screen is a list of all the object categories we are aiming to label in the image. Next to each object type is a number, which is a keyboard shortcut for that object in order to create a label of that type (this can often be very time saving once you get into the flow of things).

There are also two more options for labelling, “Good coverage?” and “Good quality?”, which we have reserved for reviewing (discussed below).

Note that some of the object you want to label may be too small to do at the current resolution of the image (depending on how sensitive your computer mouse is and the size of your screen). Labelbox (and other labelling softwares) can accommodate this by allowing you to click and drag the image around and scroll with your mouse to zoom in on particular regions of the image. For example, we may wish to zoom in on the group of individuals next to the red car in the image shown above:



To begin labelling, either press the corresponding number your keyboard or simply by click the object name from the left (e.g. "1" on your keyboard to label a person, or simply click "person"). At that point, a crosshair will appear on the screen centred on the position of your mouse:




When you click your cursor when the crosshair is up, it will define one of the corners of the bounding box, so you can use the lines to align with the object you're trying to label. For example, we may want to label the person in the white t shirt in front of the car:



Note that we align the sides of the crosshair so that the lines are tangents to the boundaries of the person (the top of their head and their elbow), as best as we can visibly determine. After clicking once to determine the first corner of the box, you simply move your cursor over until the edges of the box just skim the opposite boundaries of the object in question (in this case, the person's other arm and the **visible bottom** of their lower half).

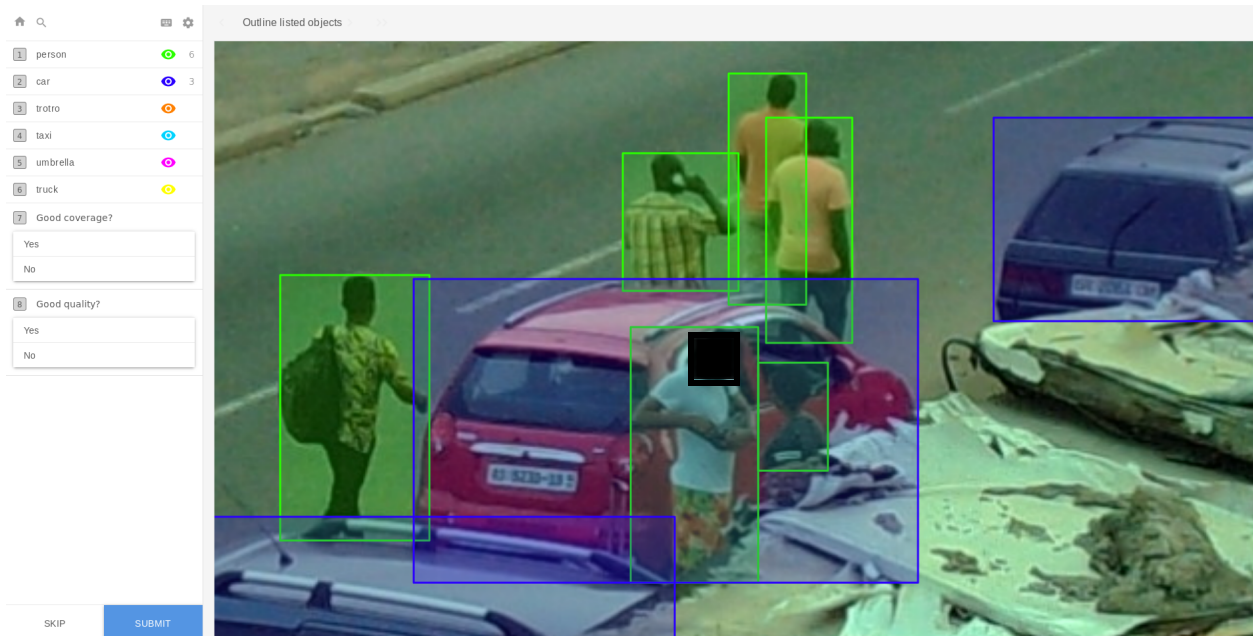


And then simply click again to finish creating the box. Once a box is made, it can also be easily amended. For example, if we accidentally made the box **too small** such that **not all visible parts of the person** are in the box, such as below:



By clicking on the white corners of the box, we can drag the corners around until they align properly with the visible edges of the person.

Then, simply label every visible object in the image, as listed above (and on the left hand side of the labelling interface on Labelbox):



For convenience, Labelbox will provide different coloured translucent boxes for different categories of objects (all people will be green, all cars blue). Note that the boxes are very often overlapping and that **even when most of the object obscured/hidden behind something else**, if we can identify the object, we still aim to put all visible parts of it in a bounding box. As you become familiar with this process, you will likely find that you can produce high quality bounding boxes in just a few seconds per box.

A **fully labelled** image might look like the following:



Once you are sure you have labelled **every easily visible instance** of **every object class** in the image, you can **click the blue “submit” button** in the lower left hand corner, **or simply press the “e” key on your keyboard** to submit your label. Note that some people find it easier to first label every object of a given class (e.g. every car), then move onto the next category (every trotro) sequentially. You will then be shown the next image.

The images you are shown are allocated **completely randomly by Labelbox**. This ensures each labeller has a fair chance of receiving images that have very few or rather a lot of objects that need labelling. For the images from some locations (e.g. University of Ghana), there may be very few or no objects that need labelling at all. In other locations (e.g. Jamestown or Nima), there may be dozens of people, cars and umbrellas. However, on average, there will probably not be many images of this sort.

Furthermore, some categories have way more instances than others and it is likely that many images will not contain objects such as cookstoves or trucks. However, we ask that you **remain vigilant** and **inspect each image for very rarely occurring objects such as loudspeakers or animals**. This will help ensure we produce a high quality data set.

In the top left corner of the labelling interface is the home icon. Clicking it will take you to an overview of your labelling so far, allowing you to keep track of how many images you have labelled.




The goal for this pilot labelling is for each labeller to label **20 images a week on average**. If you are unable to do this or have any issues, please email back to us and we can advise and assist.

Labelling Guidelines:

In order to produce high quality labels that will allow for the best possible machine learning algorithm to be developed and trained, we require labels to be created in line with the following guidelines:

- 1) **All visible parts** of the object must be included in the box, but **no more**.
- 2) When the boundary of the object is blurry, exercise your best judgement on where the box should go.
- 3) Label objects even when they are not fully visible/partially obscured.
- 4) Only apply a single bounding box per object and ensure every object has its **own** bounding box.
- 5) Do not attempt to label objects that become too difficult to discern. For example, in the upper right hand corner of the example image used above, the individual cars, people and distant umbrellas become too difficult to distinguish. You should not spend enormous amounts of time trying to label them. This requires some personal judgement but a good rule of thumb is whether or not *you* can easily and immediately see what the object is and where its boundaries are.



- 6) Use the most specific category of object available.

In particular, although the taxis and trucks in these images may also be categorised as “cars”, we require that you only use the “car” label on an object bounding box if it is **not** one of those two other categories. In other words, do not label a taxi or truck as a car. However, for person, vendor and cooking bowl, please label the “vendor”

in line with the guidelines given in the Examples document below **as well as** labeling the person with their own bounding box, and the cooking bowl separately.

We please ask all labelers to consult the examples document and look through the examples (though the document appears to be long, it is predominantly images which will only take a few seconds each to review) and familiarise themselves with the desired labeling guidelines, before labeling:

[Object Detection Labeling Examples](#)

At some point, just so we have an idea about the quality of the labels, we may get people to (anonymously if possible) review each others labeling and judge the quality and coverage of the labels we generate (simply by marking the images with yes/no categorisations of good quality/coverage). This is mostly so we have an idea of the quality of the labels (which will be imperfect to some degree) and potentially to provide feedback throughout the process, though information on this will be sent in a separate document.

As ever, please email r.nathvani@imperial.ac.uk if you have any questions about the above.

Cheers all,
Ricky and Emily