

Timelapse Image Recognition Guide

A guide to importing and using image recognition data

The screenshot displays the Timelapse software interface. The main window shows a timelapse image of a savanna with two giraffes. One giraffe is highlighted with a blue bounding box and labeled "giraffe 1(1)". Another giraffe is labeled "giraffe 59(59)". The ScoutGuard logo and timestamp "08.21.2013 17:14:36" are visible at the bottom of the image.

Overlaid on the right is a dialog box titled "Select and view a subset of your files". It contains the following information:

- What:** You may want to view only a subset of your images and videos that fit some criteria of interest to you.
- Solution:** Specify the search terms that describe your criteria.
 - Each row below reflects your data fields or (if enabled) specific recognition data.
 - Select one or more rows and adjust its values to reflect your search criteria.
 - If you have selected multiple terms in the lower area, select how those terms should be combined.
- Result:** Only those images and videos matching your search criteria will be displayed.
- Hint:** Glob expressions are case sensitive and allow wildcards as follows:
 - * matches any number of characters and ? matches any single character
 - [abc] matches one of the indicated characters; [a-z] matches one character in the range of indicated characters.

The **Image Recognition** section is checked, with "Use recognition" selected. The recognized entity is "giraffe" and the confidence is set from 0.80 to 1.00. A color bar below the confidence range shows a yellow segment. The text "22 files match your query" is displayed.

The **Select images and videos that match these terms** section contains a table with columns: Select, Label, Expression, Value, and How multiple terms are combined. The table is currently empty.

Buttons at the bottom of the dialog include "Reset to All Images", "Cancel", and "Okay".

Saul Greenberg
Greenberg Consulting Inc. / University of Calgary
saul@ucalgary.ca

Timelapse Image Recognition Guide

A guide to importing and using image recognition data¹

Don't Panic! The actual process for using image recognition is fairly simple. This guide is lengthy only because it goes into detail (perhaps excessively).

This guide explains how image recognition information is obtained from 3rd party software, and how that image recognition data can be incorporated into Timelapse to make your workflow even faster and more efficient. A [companion video](#) illustrates the workflow described here.

A few things you should know ahead of time before jumping into this guide (all presented in more detail shortly).

- **You should be familiar with Timelapse.** If not, read the Timelapse Quickstart Guide (available on the [Timelapse web site](#).)
- **Other software does the image recognition.** Timelapse helps you install and run the *EcoAssist* image recognizer app on your image set, where you can choose a species identification model that best fits your needs. After importing recognitions into Timelapse, you will see bounding boxes atop of detected entities in your images, and you can select and tag subsets of image containing particular entities.
- **Image recognition isn't perfect.** It's important that you understand where recognition works and where it can fail. We strongly suggest you read this short document and watch the accompanying video presentation.
 - » [Automated Image Recognition for Wildlife Camera Traps: Making it Work for You](#). Saul Greenberg, Research report, University of Calgary: Prism Digital Repository, 2020.
 - » [Video Presentation: Image Recognition for Camera Traps: Making it Work for You](#). Saul Greenberg. *Conference: Scaling Up Camera Trap Surveys to Inform Regional Wildlife Conservation*, Columbia Mountains Institute of Applied Ecology, May 18-20, 2020. Duration: 20 minutes.

Credits. Camera trap images used in this and other guides are used with permission or obtained from public repositories. *Sources:* Parks Canada; Lila Science - Idaho Fish and Game data set; Lana M. Ciarniello, Aklak Wildlife Consulting, Snapshot Serengeti. Everything else ©Saul Greenberg, 2022.

License terms. As the practice image and video files were provided by other agencies, their use beyond this tutorial must adhere to the license terms listed in *Description*. [LicenseTerms.Credits.pdf](#), as found in the *PracticeImageSet* folder.

Table of Contents

What is image recognition?	3
Two types of image recognition: Detection vs. Classification	3
Image recognition is not full proof	3
Why use automatic image recognition?	3
EcoAssist / Megadetector image recognizer	4
Downloading EcoAssist	4
Running EcoAssist and using its results	4
The Image Recognition Practice Set	5
How Timelapse interprets recognition data	6
» Detections are displayed as labeled bounding boxes	6
» Displaying classifications	7
The meaning of confidence levels	7
Setting bounding box preferences	8
Selecting images fitting a recognition criteria	9
A recognition workflow using detections	10
Prepare the template and load everything	11
Select and filter out empty images	11
Mark images with people in them	12
Mark images with animals in them	13
Mark any remaining images	13
Detailed classification	13
A recognition workflow using species classifications	14
Workflow variations	14
Skip animal verification	14
A rapid, less accurate method for tagging animals	15
Incorporating recognitions over time	15
Initial recognitions: run EcoAssist on the root folder	15
Subsequent recognitions: run EcoAssist on the sub-folder	15
Other image recognition features	16
<i>Custom select dialog extras</i>	16
<i>Recognitions menu extras</i>	16
Appendix 1: Recognition file format	17
How Timelapse associates recognition data with an image	17
Appendix 2: Trouble-shooting file importing	17
Appendix 3. Megadetector without EcoAssist	18

What is image recognition?

A holy grail of camera trap analysis is automated image recognition, where a computer automatically inspects, analyzes and tags your images. While we are not quite ready to let a recognizer do all the work, image recognition can help ease your workflow, as long as you understand its many subtleties.

You can experiment with image recognition in Timelapse (and we do encourage that). Still, you will get better results if you have: some knowledge of how image recognition works; its benefits and weaknesses; how Timelapse shows you image recognition results; and the various ways you can incorporate recognition data into an effective workflow.

Internally, recognizers incorporate one or more *models*, which it uses to detect and classify entities in your images. These models are created by training it on sample data, such as previously tagged images identifying species of interest in a geographic area. Some recognizers, including EcoAssist, let you choose a model that best fits your purposes.

Two types of image recognition: Detection vs. Classification

A *detector* detects whether something is in an image or not. For each suspected *detection*, a detector:

- assigns a coarse identifying label to it (e.g., *empty, animal, person, vehicle*);
- locates it via *bounding box* coordinates, which can be used to draw a rectangle on your image that identifies the suspected entity;
- assigns a *confidence value* indicating the likelihood that it is correct.

A *classifier* analyzes the sub-image contained by the detector's bounding box, where it performs a fine-grained classification chosen from a set of known categories. A classifier:

- produces a list of possible *classifications* for each detection;
- includes a *probability value* that very roughly indicates the likelihood that the classification is correct (e.g., for wildlife species the list may be elk at high probability, moose at lower probability, deer even lower, etc.).

Image recognition is not full proof

While image recognition can work reasonably well, it is not a magic bullet.

Several types of recognition errors can occur depending upon the image recognition algorithm, how its underlying recognition model was trained, how well the images you submit match that model, and the image contents.

Recognizer errors fall into these categories.

- *False positives*: it identifies an entity as being detected in an image when nothing is there.
- *False negatives*: it does not detect an entity when that entity is present.
- *Incorrect identifications*: it detects an entity, but incorrectly labels it (e.g. an animal instead of a person for a detection, or a wrong species for a classification).
- *Ambiguity*: it detects several overlapping and possibly conflicting detections.

Why use automatic image recognition?

While the ideal recognizer would simply tag all your images for you, recognition errors are a fact of life. Unless the error rate is something you can live with, you should not blindly accept all recognition results. Instead, you should consider recognitions as predictions that can augment, rather than replace, your workflow. As we will shortly see, you can use Timelapse to show only subsets of images matching a particular recognition category. You can then tag each category in bulk while still checking and correcting for occasional errors. This allows you to tag most images quickly, leaving far fewer images requiring closer inspection.

Typical benefits are listed below.

- **Eliminate images that are likely empty from any further analysis**, such as images with no detected people or animals in them or images with only very low confidence. This is especially valuable in cases where your cameras generate a large number of empty images, e.g., because you have your camera on a timelapse mode, or because wind effects frequently trigger your camera in motion-detection mode.
- **Identify images containing people**. This is especially important if your agency has a privacy policy that requires those images to be handled somewhat differently from other images, or when you have a different set of analysis criteria for people vs. animals.

- **Identify images containing vehicles.** As some camera traps are set near a roadway, vehicles passing by can trigger the camera. If these are not needed, they can be selected and eliminated from further consideration.
- **Identify images containing animals,** where you will inspect those in detail, for example, to classify or count what animals are present.
- **Focus on a particular species.** If classification data is available, use the classification categories to eliminate species of little interest, or to quickly focus on those images containing a particular classified species.

EcoAssist / Megadetector image recognizer

Timelapse does not do image recognition by itself. Rather, it relies on third party image recognizers do the image processing. In this guide, we focus on *EcoAssist*: it is integrated into Timelapse, and works very well with it.

EcoAssist serves as a front end to Microsoft's *MegaDetector*, a reliable image recognition system. *EcoAssist* encapsulates the messy details of installing and using *MegaDetector*. It also lets you choose between pre-trained species identification models for one that best for your needs. *EcoAssist*, in conjunction with *MegaDetector*, analyzes images in a folder that you select, where it generates a recognition file describing those images. You can then import that file into Timelapse and use recognitions to simplify your workflow, described shortly.

MegaDetector is a popular image recognizer for camera traps. It is specifically designed for wildlife-oriented camera trap images, and is popular with many ecologists around the world. Timelapse is normally recommended as the 'front-end' for visualizing *MegaDetector* recognition data, and many organizations have used both to efficiently tag massive numbers of images.

Note: EcoAssist species identification models. Expect *EcoAssist*'s list of these models to grow over time. If the available models don't match your needs, you can still run *EcoAssist* with no model selected. *EcoAssist* will then generate detections with broad animal/human/vehicle classifications, which is still very useful. For example, you can efficiently eliminate empty images (such as those triggered through wind effects), which can significantly reduce the number of images you need to inspect. If privacy maintenance is a concern, you can identify those images with people in it and tag them for deletion.

Downloading EcoAssist

Downloading *EcoAssist* is straight forward. It will likely take 10-20 minutes, as various packages need to be installed to do image recognition. Try not to interrupt the installation. And, you can still do other work while its installing.

1. Select *Recognizer |EcoAssist Image Recognizer | Install EcoAssist image recognizer (via browser)...* An instructional dialog box will appear. Timelapse will also tell you if *EcoAssist* is already installed.
2. Click *Okay* to continue. A browser should appear displaying a *Windows Install* web page. If it doesn't, go to addaxdatascience.com/ecoassist-windows/. After reading the instructions, press *Click here to download the installation file* on that page.
3. A command prompt window will appear that will ask you if you want to install it only for yourself or for everyone on your computer. Choose whichever option sounds right for you. Let it run! You'll see lots of text scroll by. For the most part, you can just ignore all of it until it says its done.

Note. If your computer is managed by your organization, there is a good chance that your systems people have locked it down, where it prohibits installing applications. If *EcoAssist* does not install, ask your system person to do it for you. Similarly, anti-virus, VPN, proxy servers or other protection software might interfere with the installation. If you have those enable, you may want to turn those off.

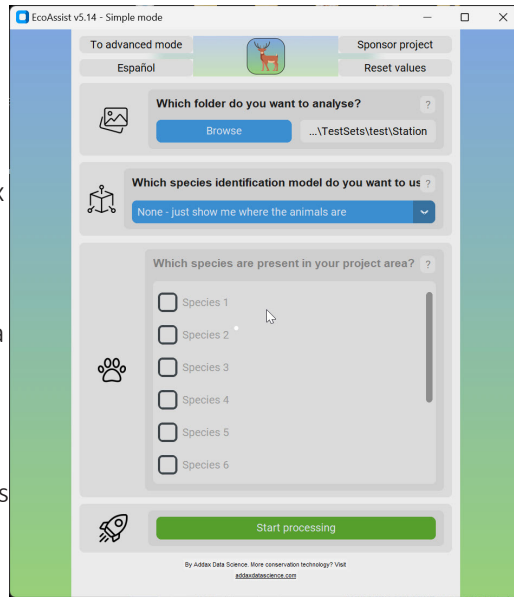
Having said that, installing *EcoAssist* on my own personal computer and laptop went flawlessly. Hopefully, you will have a similar experience.

Running EcoAssist and using its results

Running *EcoAssist* is also straight forward; these instructions are likely not even needed.

1. Open Timelapse on an image set, and select *Recognizer|EcoAssist Image Recognizer|Run EcoAssist on this folder...*
2. Select the folder containing images (possibly located in sub-folders) that you want to recognize. Usually,
 - » select the root folder of your image set if you want it to recognize all images in your image set; or
 - » select a particular sub-folder if you want it to recognize just those images (e.g., you may have already run the recognizer on your


images, where you now want to just run it on a folder containing new images that were retrieved from a camera). An instructional dialog box appears telling you that *EcoAssist* is starting up. Click *Okay*. The *EcoAssist* window will appear after a few seconds. It asks which species identification model you want to use.



- » **None:** will show detections with broad animal/personal/vehicle classifications. Choose this if the other models don't seem to fit what you are looking for. Or,
- » **Choose a model** using the drop-down menu, if it seems to fit your purposes. It then lists species known to that model. While you can select all of them, its best to select only those you know are present in your area.

Choosing a new model may initiate a further download. If so, just follow the instructions.

3. Click *Start processing*. A small progress dialog will appear. When done, *EcoAssist* will tell you where the recognition file is located (it should be in the same folder that you asked it to recognize).
4. Select *Recognizer | Import recognition data for this image set...* to load the recognition file into Timelapse. Timelapse will then display recognized entities, if any, within bounding boxes.




Note. After processing your images, *EcoAssist* generates a file called *timelapse_recognition_file.json*. This file contain recognition data for every image file found in the folder you chose in Step 2. That data also includes the image file's *relative path* and *file name*, which Timelapse uses to locate the file. Consequently, keep the recognition file in that folder. If you move it to another folder, Timelapse may not be able to find your image files as the *relative path* would not match.

Also read the *Working with JSON recognition files* at the end of this chapter. It explains how to make sure Timelapse can find the files specified in the recognition file, and how to include those recognition files in your workflow. It also describes how you can create multiple *recognition* files, each representing a sub-folder within your root folder that you can add incrementally as new images become available over time.

The Image Recognition Practice Set

Follow along with the upcoming explanations using the *ImageRecognitionPracticeSet*, downloadable from the [Timelapse web site](#). It contains a template and a single *Station1* folder containing image files from the [Snapshot Serengeti](#) project. It also contains a previously generated *.json* recognition file. This particular file contains both detections and classifications data.

ImageRecognitionPracticeSet



To use the Image recognition practice set:

- download and unzip the folder to a convenient place
- start Timelapse and load the template in the practice image set as normal;
- import the *.json* file by selecting *Recognitions | Import image recognition data for this image set..*

If you have run *EcoAssist* on your own files to produce your own *.json* recognition file, you can use that instead of the practice image set.

How Timelapse interprets recognition data

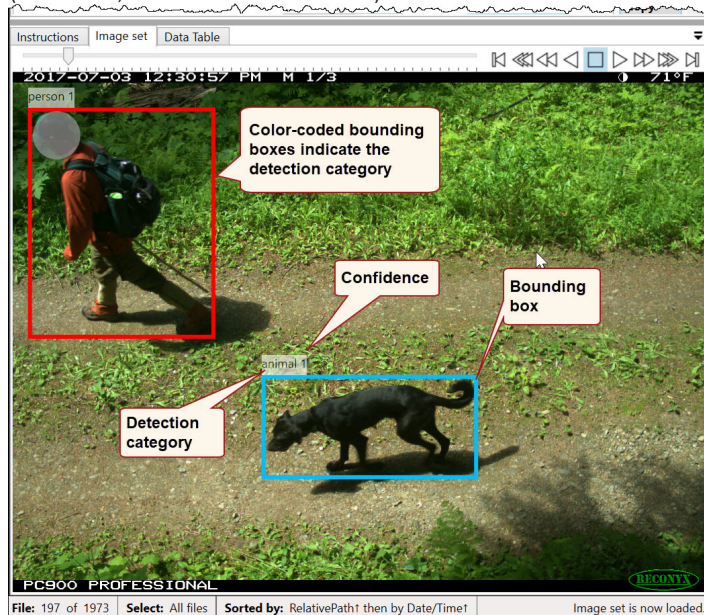
After you import the recognition file into Timelapse, you will be able to view recognitions atop your images. You will also be able to streamline your workflow by selecting particular image subsets based on these recognitions, as discussed shortly.

Detections are displayed as labeled bounding boxes

Timelapse uses the data in the recognition file to display each detected entity as a colored bounding box drawn atop each image. Each box:

- is labeled with what it thinks the entity is (e.g., person, animal);
- is colored to distinguish the detection category (e.g., red for person, blue for animal);
- includes a confidence value of correctness ranging from 1 (high confidence) down to 0 (low confidence).

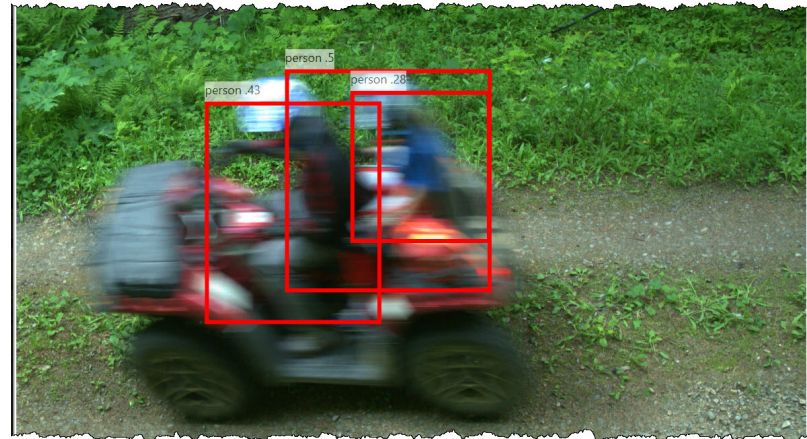
This image, for example, shows two correct high confidence detections within its labeled bounding boxes. One contains a detected person (confidence score = 1, red box). The other correctly identifies a detected animal (blue box, confidence score = 1).



Note 1. Bounding boxes can be temporarily hidden by pressing and holding the **H** key while your cursor is over the image. This is helpful if you need to inspect a portion of the image directly under the bounding box border.

Note 2. Select *Recognitions/Set bounding box options...* to set a bounding box display threshold. Timelapse will display bounding boxes only for detections above that threshold. This helps reduce clutter.

This image has detections of several people, although at lower confidence. It also produces somewhat ambiguous results as it shows various possibilities of what it thinks is a person. For example, you cannot simply count the bounding boxes, and assume that there are three people in the image.



Other errors are shown below: a vehicle labeled as an animal, an animal detected when nothing is there, and missed wildlife.



Incorrect identification



False positive



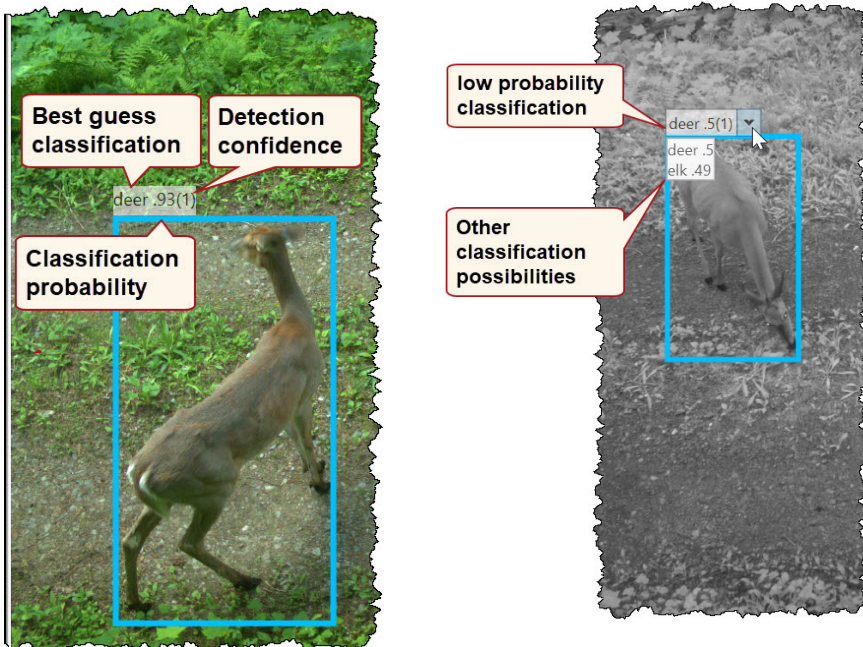
False negative

Displaying classifications

Classifications, if available, work with detections. Their visualizations are similar, except that extra information is added.

The left image correctly classifies the detected animal as a deer. The numbers that follow provide the individual confidence values as probabilities: the classification probability is .93, while the detection confidence is (1). While the 'animal' label is not included, it is implicit by its classification (a deer is an animal). The bounding box also uses the animal color.

The right image also correctly classifies the detected animal as a deer, except this time with a much lower probability value (.5). Because of its low probability, a drop down menu is included that lists alternate possibilities of lesser probabilities, in this case an elk at .49. When multiple classifications appear, the bounding box label displays the highest probability prediction. If no classifications are available, then the label only shows the recognized detection category (e.g., Person, Animal, Empty).



The meaning of confidence levels

Confidence is best seen as a very rough indicator of how likely it is that a detection or classification is correct. The catch is that the reliability of a particular confidence value can vary greatly from dataset to dataset, as explained in the technical note on the next page.

To mitigate this variability, the *json* recognition file specifies a *typical detection threshold*, and a *typical classification threshold*, where a confidence and probability value is normally interpreted relative to those thresholds.

- **At or above the threshold:** the recognition is likely mostly correct, although the occasional error may occur.
- **Somewhat below the threshold:** recognitions are somewhat more suspect, where increasingly more false positives will appear.
- **Far below the threshold:** the lower the value, the more likely the recognition is wrong.

Timelapse uses the various thresholds in the *json* file to set defaults for:

- when to display bounding boxes, where detections below a confidence threshold are hidden
- in the *Select | Custom select* dialog, what confidence ranges are best used for selecting images that contain a particular recognized entity.

Because the thresholds in the *json* recognition file are not completely reliable, you can optionally set the confidence values you want to use to ones that better fit your images, as discussed next.

Technical note

Confidence of detections vs. classifications. Each detection has a single confidence value indicating its relative likelihood that it is a correct detection. Classifications differ. As the classifier produces 1 or more predictions per detection, it assigns each classification with a relative probability of being correct, all summing to 1. The classification probability is used as a rough indicator of confidence, i.e., how likely it is that the entity was correctly recognized.

Confidence values and 'mostly correct' thresholds. Different versions of MegaDetector use different strategies for generating confidence values. As such, the confidence threshold above which detections and classifications are 'mostly correct' can differ greatly, as they depend upon the MegaDetector version used. For example,

	MegaDetector version 4	MegaDetector version5
Typical detection threshold	0.8	0.2
Typical classification threshold	0.05	0.75

As these values can change across future versions of MegaDetector, version-specific values are included in the .json file. Timelapse then uses those values to determine the default confidence levels in the *Select* dialog.


As these default values are (at best) an approximation, the best way to judge a reasonable detection or classification confidence level is by reviewing your images to get a sense of how well recognitions perform at different levels.

Bounding box display threshold. In the *Recognitions/Set bounding box options...* dialog(see next section), the initial default value for displaying bounding boxes is calculated using the detection confidence threshold specified in the image recognition .json file. As this value can vary considerably, we suggest that any adjustments of the bounding box display threshold should be relative to this default value. Each image set remembers this preference, where it is stored in your .ddb file.

Setting bounding box preferences

You can customize how and when your bounding box appears in the preferences dialog (or just accept the reasonable defaults).

- select *Recognitions/Set bounding box options...* to set a bounding box display threshold;
- Click the *Annotate bounding box* checkbox to either include or exclude the recognition labels atop the bounding box.
- If you have color blindness issues, the *Use color-blind friendly colors* will adjust the bounding box colors accordingly.
- Use the slider to adjust the confidence threshold, where bounding boxes are displayed only for recognitions whose confidence is at or above that threshold. In the example, below, only detections with a confidence of .5 or higher will be displayed. This is discussed in more detail in the technical note on the next page.
- Pressing *Reset* restores these settings to their default values.



Experiment with the *Always display bounding box* setting by reviewing your images. While it is tempting to set the threshold to a low value, it could add considerable visual clutter by showing bounding boxes around incorrect recognitions. Alternately, setting it too high will mean that some correctly detected entities will not show their bounding boxes. The best value will likely be at, or somewhat below or above the default threshold.

Workflow tip. As discussed next, you can use *Custom Select* to selectively view as subset of images that include detections and classifications at a given confidence range. When you selected images using a confidence range lower than the threshold set in the *Preference* dialog, Timelapse will display bounding boxes down to that lower confidence range (because you are interested in those images), regardless of the *Preference* dialog setting.

Selecting images fitting a recognition criteria

As discussed in other Timelapse guides, *Select | Custom Selection* displays a dialog box that allows you to select a subset of your images to view according to some criteria.

When a recognition file has been imported, that dialog box will now include an *Image Recognition* area, illustrated and annotated below. (Confidence defaults are from MegaDetector version 4).

? Select and View a Subset of your Files

What: You may want to view only a subset of your images and videos that fit some criteria of interest.

Solution: Specify the search terms that describe your criteria.

- Each row below reflects your data fields or (if enabled) specific recognition data.
- Select one or more rows and adjust its values to reflect your search criteria.

Result: Videos matching your search criteria.

Hint: Use sensitive and allow wildcard characters and ? matches any character; [a-z] matches one of the indicated characters; [a-z] in the range of indicated characters.

Image Recognition

Use recognition Rank by confidence Show only those files the recognizer did not process

Recognized entity: animal

Confidence: from 0.80 to 1.00

Include all files in an episode when at least one file matches

270 files match your query

Select images and videos that match these terms

Select	Label	Value	How multiple terms are combined
<input type="checkbox"/>	File		These terms are combined using AND: returned files match all selected conditions.
<input type="checkbox"/>	RelativePath folder includes subfolders	=	
<input type="checkbox"/>	DateTime	≥ 25-Jun-2017 06:52:16	
<input type="checkbox"/>	DateTime	≤ 25-Jun-2017 06:52:16	
<input type="checkbox"/>	ImageQuality	=	
<input type="checkbox"/>	Delete?	=	

Reset to All Images

270 files match your query

Cancel Okay

The various recognition features of this dialog are detailed below.

- **Use Recognition** checkbox, when checked, will enable the recognition controls, which will let you set recognition criteria for selecting files
 - **Recognized entity** contains a drop-down menu that lets you select files by detection or classification category. Detections are always listed at the top.
 - » *All* – the image includes a detection of one or more animals or people within the given confidence range. All excludes images where the image recognizer has not identified any detections, even if the confidence is set to 0.
 - » *Empty* – allows you to rapidly select and inspect images where the image recognizer has not identified any detections, or that contain only low-confidence detections.
 - » *Person*: the image includes a detection of one or more people within the given confidence range
 - » *Animal*: the image includes a detection of one or more animals within the given confidence range
- Classifications are only listed if they are available, and can include a long list of species.

Note. The detection and classification categories are specified in the recognition file produced by the *MegaDetector*. The ones used above are an example. Your own categories may differ from these.

- **Confidence Range** is adjusted by either a range slider or the up-down boxes. It allows you to specify a recognition confidence range of interest for each entity type. As discussed in the technical note, its default values are calculated from the typical detection and classification thresholds, .
- **Rank by confidence.** When checked, all images of the selected entity type will still be returned, but sorted by confidence value. That is, the first images will be the ones where the recognizer has the highest confidence of correctness, followed by images with progressively lower confidence of correctness. For images with more than one detection, the highest detection or classification confidence value is used. If the selected entity is a detection, the detection confidence is used. If the selected entity is a classification, the classification probability is used.

The examples below illustrate working with confidence. Understanding the subtleties – especially for *Empty* – will be helpful to using recognition data effectively. This example illustrates confidence defaults based on MegaDetector version 4, which uses a typical detection threshold of 0.8 and a typical classification threshold of .75. You will need to calibrate these ranges to the default that actually appears when you use Timelapse.

- **Entity = Empty, confidence range is 1 - 1** selects all images where the image recognizer has not identified any detections. You may use this to rapidly inspect (and occasionally correct) images with nothing in them.
- **Entity = Empty, confidence range is 0.8 - 1** selects all images where the image recognizer has not identified any detections, and all images with 0.2 or less confidence (i.e., $1.0 - 0.8 = 0.2$). While similar to the above, it should include more images that likely don't have anything in them (i.e., erroneous detections). A few more correct detections may also appear.
- **Entity = Empty, confidence range is 0.8 – 0.99** selects all images with 0.2 or less confidence (i.e., $1.0 - 0.2 = 0.8$) but excludes images where the image recognizer has not identified any detections. You would use this a part of a workflow where you first wanted to go through images with no detections (1 – 1), and then separately go through images with poor detections.
- **Entity = animal, confidence range is 0.8 – 1.0** displays images containing at least one detected animal at 0.8 confidence or above (i.e. high confidence animal images that you may want to rapidly inspect for correctness). You may use this to rapidly inspect images to ensure animals are in it, and correct the occasional error.
- **Entity = all, confidence range is 0.5 – 0.8** displays images where at least one detected entity is between 0.5 – 0.8 confidence range. You may use this to rapidly inspect detections that are more error-prone.
- **Entity = all, confidence range is 0.1 – 0.5** displays images containing at least one detected entity between 0.1 – 0.5 confidence, (i.e., images containing detections where none of them exceed 0.5 confidence. You may use this to rapidly inspect and eliminate detections that are likely wrong.
- **Entity = Black bear, confidence range is 0.8 – 1** selects all high probability black bear classifications. You would use this a part of a workflow to perhaps check and quickly label those images as a bear.
- **Rank by confidence** is useful in your workflow for determining – somewhat qualitatively – the confidence range where that recognizer

seems to work well vs. poorly. Since files are now sorted by confidence, you can play the images from the beginning, and scan them to check if they have the correct entity. When errors increasingly appear, you can check the confidence value (displayed on the bounding box) to see where the recognizer no longer provides generally useful recognition results. You can then set the confidence sliders to select the mostly correct range.

A recognition workflow using detections

Note. The following workflow description uses the image recognition practice set which contains only 272 files. It also used confidence values based on MegaDetector version 4, so the confidence values you use in other MegaDetector versions may differ. In practice, the suggested workflow reaps substantial benefits when used over larger image sets.

If you have your own image set and *.json* file, you can mimic these steps.

Recognition workflow strategies depend upon you:

- selecting a subset of your images to display based on the detector's predictions (via *Select / Custom Selection*),
- bulk- tagging that sub-set in a single action
- rapidly reviewing those images for tag accuracy and for correcting errors.

While it may seem like the following steps adds work, you will find that – depending on your images – you can quickly reduce the number of images you need to manually inspect and tag to a small subset of the original files.

The workflow described below is based on using detection information only. This particular workflow assumes that there is a very low tolerance for errors, and is perhaps the most thorough one you can do. It also uses the typical detection and recognition defaults, which can differ from your own images.

You will inspect and tag: all empty images, images with people in it, and images with animals in it. Afterwards, you can classify just the animal images.

Workflow Tip: This workflow is just one of several possible. The one you choose depends on the nuances of your work. A subsequent section will describe several other workflows that may better fit different needs. We fully expect you to create your own workflow protocol that fits your own needs – these descriptions are included just to give you a flavour of the possibilities. Still, go through this one as it described the details that should be considered in all workflows.


Prepare the template and load everything

1. Add the following fields to the Timelapse Template: three flags called *Person*, *Animal*, and *Empty*. These are already including in the image recognition practice set. Note that these labels have no special meaning to Timelapse. Their role in this example workflow is to track and bulk-tag what detections have been accepted and corrected.
2. Check that the appropriate *.json* detection file is in the same folder as your template.
3. Load in your image set into Timelapse as normal.
4. Select *File|Import image recognition data for this image set* to import the *.json* file. After this step, you should see blue or red bounding boxes in images, each box indicating where the detector has spotted a person or animal.

Select and filter out empty images

Your goal in this step is to remove empty images (i.e., where the detector has not detected anything in it) from any further consideration, while correcting any detection errors.

5. Use *Custom Select* to select the images where the detector has not detected any entities.

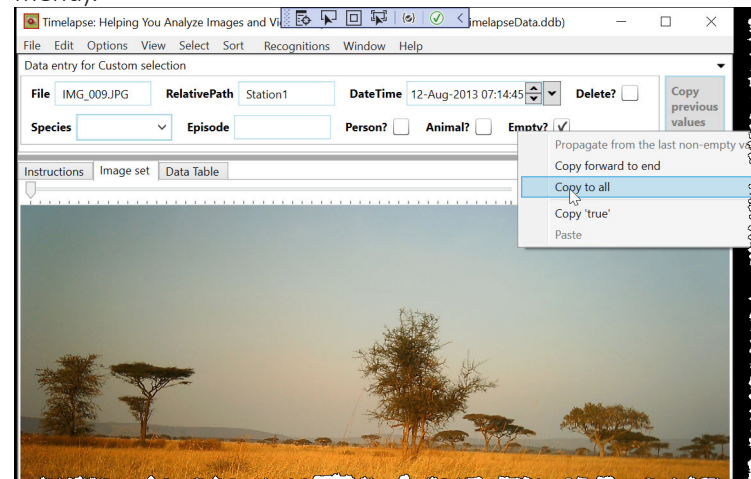


- a. Click *Use Recognition*
- b. Set *Recognized entity* to *Empty* and a confidence range from 1 to 1.
- c. In the image recognition practice set, this will match 113 out of the 272 total files. Click *Okay*.

Note. The number at the bottom right above indicates how many images match that selection. The same number is repeated at the bottom of the dialog box.

Timelapse is now displaying the 113 images that the detector has considered *Empty* at a high confidence. While we would expect most of these to be empty, they have to be checked.

6. Tentatively tag these images as empty by setting all of their 'Empty' fields to true.
 - a. Checkmark the *Empty* flag of the current image.
 - b. Select *Copy to all* on the *Empty* flag (right-click raises its context menu).



7. Starting from the 1st file, review the images until you are satisfied that the detector was reasonably accurate at identifying images without animals or people in it. You can do this efficiently by using one or more of these strategies:
 - » use the *FilePlayer* to automatically play the images at various speeds. *Options|Adjust File Player Speeds* controls playback speed.
 - » use the *Overview* to display multiple images at a time, where you can also use the *FilePlayer* control to scroll through pages of images;
 - » use the arrow keys to rapidly move through images.

Reminder. A custom selection will show the last image (or as close to it as it could get) that was previously selected. Remember to scroll back to the first image before reviewing them, as otherwise you will miss some of them.

8. Correct erroneous *Empty* detections as needed. If the detector is working well, the vast majority of images you see will be empty. If you do see images that contain animals or people (false negatives), you can uncheck

the *Empty* field for those images. At the same time, you can broadly classify those by checking their *Person* or *Animal* field and (if you want) set its species. For example, IMG_148-149 has an unidentifiable animal in it (very small, on the middle far left), and IMG_261-263 has a zebra (you may have to use the *Gamma* setting on the *Options | Temporarily adjust image appearance* to see it in the dark).

Workflow tip. Create two *QuickPaste* entries to do the above with a single mouse click. The first clears *Empty* and sets *Person*, while the second clears *Empty* and sets *Animal*.

9. Depending on your tolerance for errors, you may want to review all images carefully and correct for false negatives, or just look for glaring errors and correct those. Alternately, you may want to stop reviewing empty images at some point, if you feel the detector's accuracy is good enough for your needs.


At this point, we have identified 109 empty images, and corrected 4 errors in just a few minutes. 40% of the total files have now been processed!

Mark images with people in them

Some organizations require, usually due to privacy policies, that images with people in it be identified and even discarded. Alternately, you may be counting people in your images, or you may not be interested in images with people in the context a wildlife survey. In any of these scenarios, it's useful to rapidly identify the subset of images that contain people.

Your goal in this step is to identify as many images as possible with people in them. The process is somewhat similar to how we deal with empty images.

10. Select the images where the detector has a high confidence that a person is in them, ranging from .8- 1, as illustrated below.



Note. As mentioned, the image recognition practice set does not contain any images with people in it. The detector has likewise not detected any people (if reducing the confidence, there will still be 0 matches). However, to show how the workflow works, the explanation will continue as if there were matches.

11. Similar to how you had previously tagged and verified empty images, tentatively indicate a person is present in all these images by setting the *Person* flag to true on one image, and then use *Copy to all* to set it for all images in the current selection.
12. Review this subset (remember to start at the first file). If you see an image that does not contain a person (false positive), clear its *Person* flag. If the image is empty or has an animal, set its *Empty* or *Animal* flag to indicate that. Again, depending on your tolerance for errors, you may want to review all images carefully and correct for false positives, or just look for glaring errors and correct those. Or, you may want to stop reviewing the images if you feel the detector's accuracy was good enough. As you've previously set them all to *Person*, you don't have to do anything more.
13. Repeat the above, but this time by setting the confidence from .8 to .9, then between .7 and .8, and so on down to 0.
 - » When you begin to see that the detector is getting things more wrong than right, don't set the *Person* flag via *Copy to All*. Rather, review the images and click only on the ones with a person in it.
 - » At some point, you can stop reviewing images below a certain confidence threshold if there are many errors.


Workflow tip. Move the *from* confidence of the range slider slowly, and monitor how many files match that confidence range. Use that information to decide how you should set the confidence ranges. For example, if only a few images match a particular range, you may want to expand the range to generate more images.

Mark images with animals in them

Your goal in this step is to identify as many images as possible with animals in them for later classification. It is identical to the way you tagged images with people in it.




14. Select the images where the detector has a high confidence that an animal is in them, ranging from .8 - 1, as illustrated below. 152 files (of the 272 total files) are in this range. Click Okay.



15. As before, tentatively set the **Animal?** flag for all those images. Review them and alter the flags as needed.

Some things to note:

- » None are empty, although a few have animals that are fairly obscure and hard to identify.
- » Some of them show bounding boxes on animals that could have been easily missed otherwise. IMG_195, 214, for example, detected animals in dark areas, where the human eye would not see them unless we investigated it with the Gamma correction facility.
- » Another example that would have been easily missed without the bounding boxes is IMG_277: a far-away animal obscured in the underbrush.



16. Repeat the process, but this time setting the upper confidence to .8, and then drag the lower confidence slider downwards. When it is at its lowest limit, we see only 5 matching files. Instead of using Copy to all, just go through them and tag whether they are **Empty**, **Animal**, or **Person** files.

Workflow tip. You could have started this process by selecting Animal, and then checking the **Rank by confidence** checkbox. The selected files will then be sorted from the highest to lowest confidence. Scrolling through them quickly then gives a sense of where the detector doesn't do as well, which would hint at what confidence intervals would give the

best results. The benefits of this are clearer when used in larger image sets.

Mark any remaining images

17. To see if you missed any images, check the **Person**, **Animals**, and **Empty** checkboxes with values unchecked. If the count is not 0, manually tag them. For example, you should find two files that have not been tagged (these two files were not processed by MegaDetector, and would have appeared if you had selected **Show only those files the recognizer did not process**).

Note: You can double check your tagged images at any time by selecting the **Empty**, **Person** or **Animal** checkbox only.

Detailed classification

At this point, you have isolated the **Empty**, **Person** and **Animal** images.

18. To classify the **Animal** images, use **Custom select**, and check **Animals** = unchecked. All the tagged Animal images (which could be a considerable number) will now be filtered from view. You can now go through the standard manual tagging mechanism to set the species for the remaining untagged images.

Workflow tip. Using the above workflow can significantly reduce the number of images requiring close inspection. In practice, many image sets can have a massive number of empty images in them: 90% empty is not unusual. Similarly, the above workflow can quickly separate images with people in them from other images.

19. To discard empty images (if desired)

- » Select the **Empty** checkbox
- » In your data entry area, check the **Delete?** flag
- » Right-click that flag to raise the context menu, and select **Copy to all**.
- » Use the **Delete** menu to delete all those files (and optionally its data).

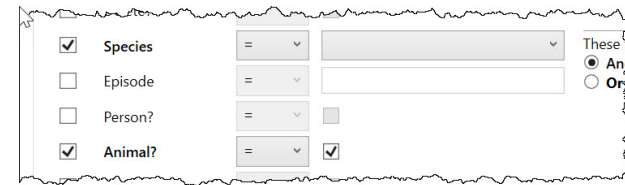
20. To discard People images (if desired), repeat step 20 expect with **Person** selected.

A workflow using species classifications

If the recognizer produces both detections and classifications (as with the image recognition practice set), further efficiencies can be achieved. The process is very similar to what was described in the previous section.

1. Follow the same steps above for:
 - a. preparing the template and loading everything;
 - b. selecting and filtering out empty images;
 - c. marking images with people in them;
 - d. marking images with animals in them;
 - e. marking any remaining images.
2. Instead of manually tagging animal species (step 19 above), you can use the classifier to bulk-tag frequently occurring species.
 - a. In the *Custom select*, choose *Animals*, and click *Rank by confidence*. Scan the returned images quickly to get a sense of which species are most frequently captured. For example, you should see many images of Thomson Gazelles, then perhaps Giraffes, followed by a small number of other species.
 - b. Uncheck *Rank by confidence* and select *gazelleThomsons* with a high confidence range. Repeat the *Copy to all* process to tentatively tag the *Species* and *Animal* data fields. Verify / correct those as before (reminder: create Quickstart entries to do this in a single click).
 - c. Repeat the above at lower confidence ranges, but only while classification errors remain fairly low.
 - d. Repeat the above process for other frequently seen species (e.g., *Giraffe*).
 - e. At this point, you will have already labeled many images with your most frequently seen species. The remaining images should be a much smaller set of infrequently seen species, or which were under the confidence limit you specified.

When you start getting to rarer species, revert to manual tagging. Turn off *Use recognition*, and then set the following:



This will now select all remaining untagged species (i.e., *Animals*, but with *Species* not set). Tag the *Species* field manually.

Workflow tip. The idea of using classifications is to bulk-tag as many images as possible. Using high-probability classifications means reviewing and correcting errors should also be fairly fast. What remains should be a much smaller set of images containing rarer or harder to classify species.

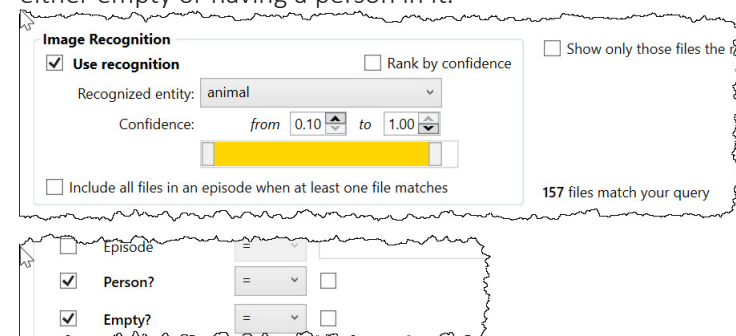
Workflow variations

The above workflow is the most thorough. However, depending on your needs, you may choose to follow one of the simpler workflows below.

Skip animal verification

You can skip the step for tagging and verifying animal detections, where you instead start classifying species immediately in the *Species* field. The difference is that you will handle mis-detected animals as you do your detailed classification, rather than as a separate step. To do this,

1. When creating the template, don't bother including an *Animals* data field.
2. Follow the workflow above for verifying and correcting detected *Empty* and *People* images.
3. To classify wildlife, set the following parameters. This will return all files with detected animals, and exclude files you've already tagged as being either empty or having a person in it.



4. Start classifying these animal images by setting its *Species* field, where you correct the occasional image that are either Empty or Animal).
5. Alternately, you can rapidly inspect and eliminate many of the animal false positives by first selecting images where the detected entity is animal, and the confidence range is (say) between 0 – 0.5. Then continue with the previous step but with the confidence range between 0.5 – 1.0.

A rapid, less accurate method for tagging animals

This approach assumes you can tolerate some errors. Note that the animal/people/empty tags in the template are required only if you want to tag them.

1. Quickly check the detector's *Person* and *Empty* results. Determine if the error rate is tolerable.

Workflow tip. You can check if the error rate is acceptable by quickly selecting the Empty and Person images, and scanning them for accuracy. For a more rigorous approach, follow this with *Select | Create a random sample from the current selection*. If you choose (say) 100 as your random sample, you can count the number of errors in that sample and infer that the entire set has that error rate.

2. If the error rate is tolerable, focus only on files with detected animals.
 - a. *Custom select* the recognized entity as *Animal*.
 - b. Slide the confidence downwards until the number of matching images does not increase significantly.
 - c. Start classifying the species of the returned images immediately.
 - d. if the returned images include some without animals (false positives), which you can just skip over when you come across them.

The above method trades off accuracy with tagging speed. There will be missed tags in this method. Some of the images detected as *Empty* or *Person* may have animals in it, as are some of the low confidence *Animal* images you have ignored.

Variation. If you simply want to ignore images with people in it vs. tagging them, don't bother creating a 'person' flag in your template. Just use the empty field instead.

Incorporating recognitions over time

In practice, most agencies add images incrementally over time, usually because new cameras have been added and/or images from SD cards have been retrieved. This section offers a workflow that efficiently incorporates recognitions of these incrementally added images.

Initial recognitions: run EcoAssist on the root folder

When you first collect images and load them into Timelapse, you can run *EcoAssist* on your image set's root folder. This will generate a recognition file for all your initial images, which you can load into WildlifeProject : Timelapse.


For example, the Wildlife project on the right initially contained images from a single camera, stored in Site01. EcoAssist was run in the root file and generate a recognition file in that folder.



Subsequent recognitions: run EcoAssist on the sub-folder

As you add folders containing new images over time , you can run *EcoAssist* on just folders.

For example, consider this update to the WildlifeProject folder. The analyst has recently installed a new camera, and retrieved its images from its SD card. She created a new folder (Site02) to hold those images, where she added them to Timelapse using Timelapse's *File | Add image and video files to this image set*.



The question is: how can Site02 recognitions be created by EcoAssist so they can be imported by Timelapse?

Approach 1. Rerun EcoAssist on all files. The simplest is to simply rerun *EcoAssist* in the root folder on all files. That recognition file will then include recognitions for all old and new files. When Timelapse reads in the recognition file, it will add recognitions for the images in the new Site2 folder, and will update the recognitions already recorded for images in the Site1 folder.

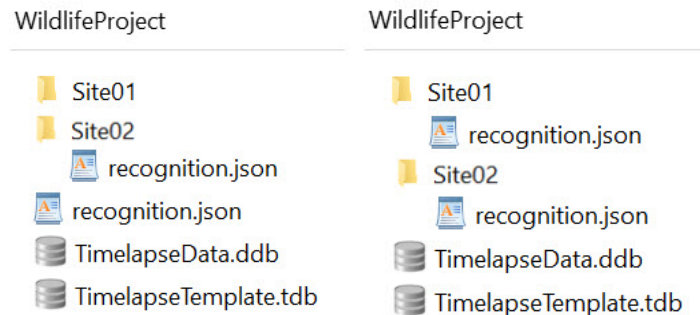
However, this approach is very inefficient, as it means you are re-running the recognizer for images that already have known recognition data. This can

add considerable time to the recognition task.

Approach 2. Run EcoAssist in the just-added sub-folder. The more efficient way is to run *EcoAssist* in the new Site2 sub-folder. *EcoAssist* will analyze only the Site2 images, and will generate a new recognition file in the Site2 sub-folder. You can then import Site2's recognition file, where, Timelapse will merge the new recognitions into its existing set of images.

The left image illustrates how a new recognition file was created in the Site02 folder and then imported. This will work, but it does leave some organizational questions about which folders that original recognition.json file (located in the root folder) are covered. This could be handled, of course, by renaming the file (e.g., recognition-Site1.json).

The right image illustrates a better organizational approach. Here, the user has recognized that folders will be added periodically, and thus creates the recognition file within each sub-folder.



In summary:

- *If your image set is complete* where you don't expect to add any images to it over time, run *EcoAssist* on the root folder.
- *If you plan to add images to your image set over time*, try to organize each incremental addition as its own sub-folder. Run *EcoAssist* in that sub-folder, where you can then incrementally add those images import that json file into your image set.
- *Don't move the json file.* Timelapse requires that the json file is located in the folder where recognitions were done, as otherwise it won't be able to locate the images associated with its recognitions.

Other image recognition features

Custom select dialog extras

Include all files in an episode when at least one file matches checkbox. When checked, all files in an episode will be selecting even if only one (or more) files in that episode match the recognition criteria. This is feature is useful for cases when the analyst wishes to review all files in an episode in order to best identify, validate the recognition prediction (e.g., some images may only show only a small, difficult to identify part of the animal as it enters the scene), or to uniformly tag the images in that episode.

For this to work properly:

- a note field must have been previously populated with episode data (using *Edit | Populate a field with Episode data...*);
- when populating, you set the episode data format to include the episode number (e.g., 24:1|7) - this is the default setting;
- populating was applied to all your images, as the episode number needs to be unique.

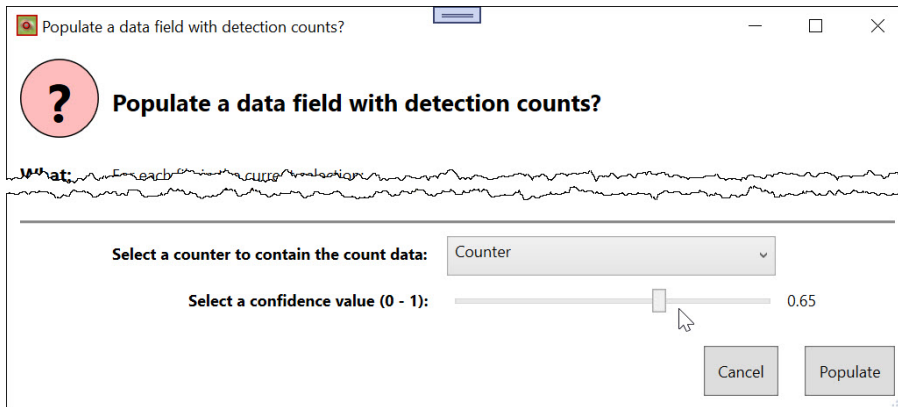
Show only those files the recognizer did not process checkbox. When checked, the custom select will return all files with absolutely no detection data, i.e., files that were not included in the .json file. If any exist, they will have to be tagged manually.

Recognitions menu extras

Recognitions | Populate a data field with detection counts... This facility, shown below, applies to all your currently selected files. It counts all detections in an image at or above a particular confidence value, and places that count in a counter data field of your choosing.

The number of detections provides a rough count of how many entities are in each image. It is an approximation, and is very much affected by the chosen confidence value (low values likely over-estimate the count, while high values likely under-estimate it).

- An approximate count for certain species may suffice. To avoid manual counting of (say) unimportant cow images, custom select on the 'Cows' classification. Then use the count facility on the returned cow images.
- You can always inspect your images and alter that count if needed.



Appendix 1: Recognition file format

If you use *EcoAssist*, you don't really need to know anything about the Json recognition files, as that file should be in the correct image folder location. However, if you use *MegaDetector* directly (see Appendix 3), the information below may help you fix problems that can occur due to your recognition file not being able to find your images.

Timelapse and recognition files are similar in that both fundamentally associate data with images found in folders and sub-folders. For the two to work together, the file paths describing an image location must match.

How Timelapse associates recognition data with an image

As Timelapse reads in the json file, it compares the file path location of the image (recorded in each json image 'file' field) with the file path location it knows about (see technical note). If there is a match, Timelapse adds that file's recognition data to the database, and associates it with the corresponding Timelapse file image. If there is no match, it skips that file.

Recall that Timelapse records its file location relative to the root folder (i.e., the folder that contains the template file). For example, consider the WildlifeProject folder on the right, which is a root folder. If the Site01 folder contained the IMG001.jpg file, Timelapse would have recorded its location as Site01/IMG001.jpg. The path indicates the location of that image relative to (i.e., within) the folder containing the TimelapseTemplate.tdb file. It does not record the 'WildlifeProject' name, as it doesn't need that information. This allows Timelapse to find the files contained within the WildlifeProject folder even if that file is moved around.

Technical Note. The Json file structure

It is helpful (although not necessary) to know what a json file normally contains. The image below portrays a stylized example. The file:

- contains some internal information about the recognizer,
- defines the detection and classification categories associated with a recognition,
- provides a list of image files, where
- each image file includes its location as a *relative path* name, and its associated detection and classification information, if any.

```

info: various information fields provided by the recognizer
detection_categories: defines several high level categories of what the detected item could be
  1: animal,
  2: person,
  3: vehicle ...
classification_categories: defines more precise categories associated with a particular classification
  0: bear,
  1: bird,
  2: canid,
  3: ...
images:
  file: Site01/IMG001.jpg the file's location
  max_detection_conf: 0.9 maximum confidence across all detections (if any) associated with this file
  detections:
    category: 1 this detection's category (animal)
    conf: 0.9, this detection's confidence
    bbox: 0.3,0.5,0.1,0.4 bounding box coordinates for the detection
    classifications: this detection's classifications (if any)
      1,0.9 this classification's category (bird) & confidence value
      more classifications for this detection, if any
      more detections for this image, if any
  more images...

```

Appendix 2: Trouble-shooting file importing

Import problems are unlikely if you use EcoAssist, as it creates recognition file in the correct folder. However, if you try to import a .json recognition file, and get a warning that no recognition data was found, the problem is almost certainly due to a mismatch between how files and paths are named in the .json file vs. how they are named in the Timelapse database. This is likely due to your recognition file being moved, or your folders being renamed before the recognition file is loaded into Timelapse

1. As explained earlier in this document, Timelapse locates files by its relative path to the folder holding the template. For example, if the

template is in the folder *MyImages*, and *IMG_01.JPG* is located in a subfolder named *Station1*, the image location is recorded relative to the *MyImages* folder:

- » its *RelativePath* is *Station1*
 - » its *File* name is *IMG_01.JPG*
 - » so its location within Timelapse becomes *Station1/IMG_01.JPG*.
2. In the *json* file, each file has an entry that describe that file's location, which must exactly match the above location. The only exception is a *json* located in a sub-folder whose paths are relative to that sub-folder, which Timelapse recognizes as a special case.
 3. To check and possibly repair path errors, make a copy of the recognition. *json* file. Open that copy in a text editor.
 - » the free [EditPad Lite](#) editor works really well for this, as the usual Windows-based editors are poor at handing really large files.
 4. Check if the paths are what would be expected. For each file, the path is recorded in the "*file*" field. For example, a correct path to the above file would be recorded in the line:

```
"file": "Station1/IMG_01.JPG".
```

If you are using MegaDetector by itself, the prefix to the path is likely either added or missing due to how files were submitted to it. For example, an incorrect path to the above file (which in this case includes the added root folder name), would be

```
"file": "MyImages/Station1/IMG_01.JPG"
```

If a prefix is missing from all paths because MegaDetector was run in a sub-folder (e.g., within *Station1*), moving that *json* into the sub-folder and then importing it would work, as Timelapse recognizes this as a special case.

```
"file": "IMG_01.JPG"
```

5. Do a global replace to add or remove the appropriate prefix, for example, to remove *MyImages/* from all files paths. Then try to import the *.json* file back into Timelapse.
6. Alternately, MegaDetector provides a [web page and downloadable application](#) for adjusting paths, as it recognizes this as a common problem. Its a bit complex to use, but it does handle complex situations, such as when you split your folders up to contain multiple Timelapse database.

Appendix 3. MegaDetector without EcoAssist

If for some reason you don't want to use *EcoAssist*, you can install the *MegaDetector* recognition software directly on your machine and generate the recognition file yourself (a bit more complex). See the Timelapse website: [Downloading the MegaDetector Image Recognition System](#).

Alternately, you can submit your images to Dan Morris (contact him at cameratraps@lila.science), who will run *MegaDetector* for you. Dan is incredibly helpful. He will:

- check that *MegaDetector* is appropriate for what you are trying to do;
- tell you how to submit your camera trap images, either through the cloud (preferred) or by sending a hard drive;
- once submitted, will have *MegaDetector* analyze your images and will return a data file (suffixed with *.json*) containing the image recognition data, which you can then load into Timelapse.

Notes

1. The only reason for not using *EcoAssist* to do the heavy lifting is if you are knowledgeable about MegaDetector and want to tweak its recognition parameters, or use it to develop your own model.
2. Several versions of MegaDetector are currently available. One is maintained by Dan Morris email:cameratraps@lila.science. Dan has actively supported MegaDetector and its community for a long time). Another version is maintained by Microsoft [AI for Good](#) research group.
3. We do expect other companies and agencies to develop their own image recognition systems, where they will produce Timelapse-compatible recognition files.