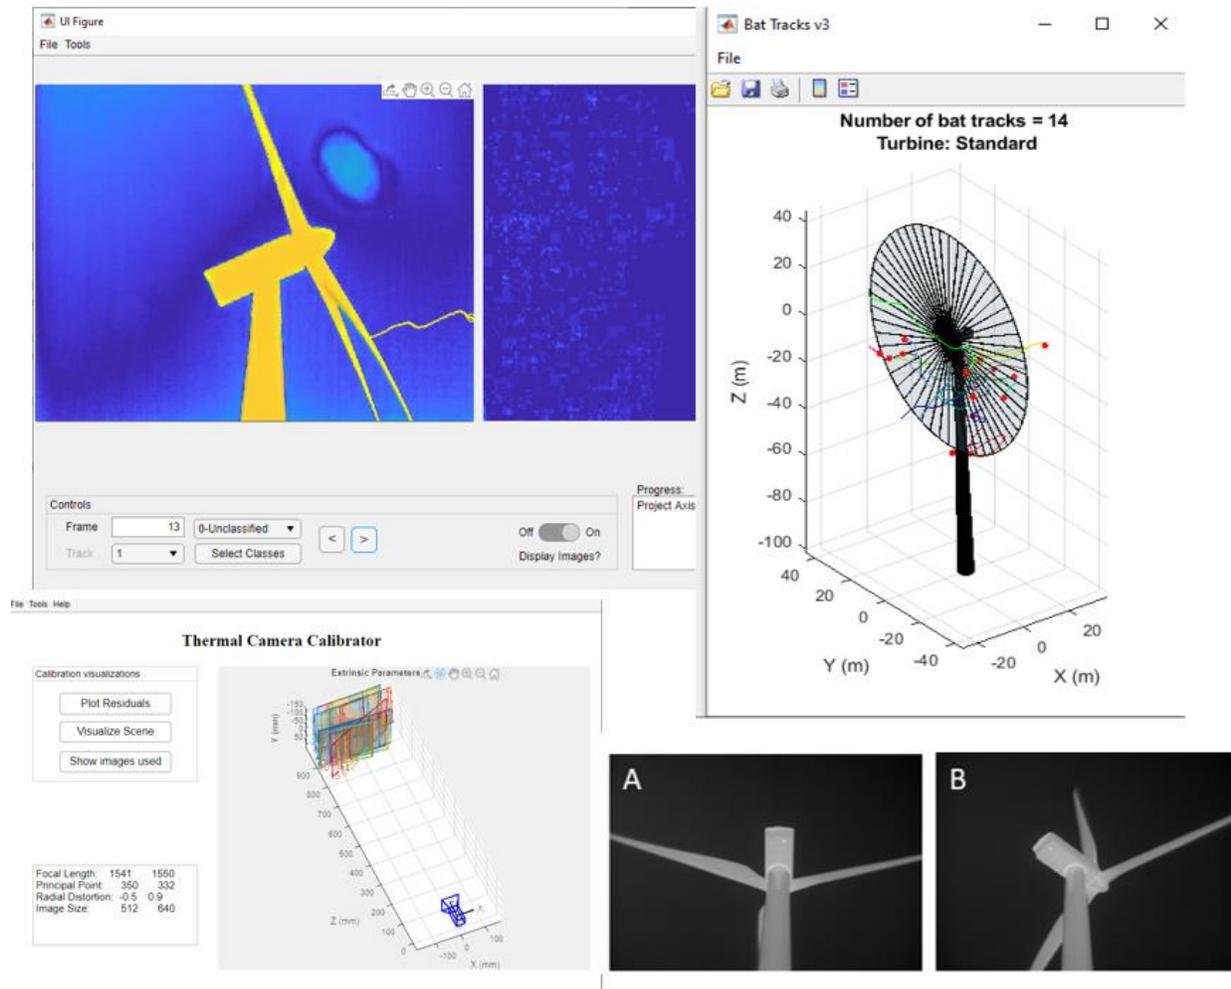# ThruTracker User Manual

Verson 1.9

Revised May 5, 2021

Aaron J. Corcoran and Tyson L. Hedrick

Contact: acorcora@uccs.edu

Updates: www.sonarjamming.com/thrutracker

Table of Contents

1. Overview

1.1 Capabilities and Limitations

ThruTracker is a powerful software platform for automated two-dimensional (2-D) and three-dimensional (3-D) video-based animal tracking. ThruTracker detects moving objects within fixed (non-moving) field of view cameras using robust algorithms and flexible settings (Chapter 5). Detections are linked together across frames to generate tracks, which can be classified into different categories such as bird, bat, insect, or counted in various ways. Those interested in 2-D tracking can skip straight to chapter 5, as chapters 1-4 pertain only to 3-D tracking.

ThruTracker uses background subtraction to detect moving objects within a stationary background. This requires contrast between the focal animal and the background. There are several requirements for generating 3-D tracks. Specifically, two or more cameras must have an accurate time synchronization. Cameras must also be fixed in space and not move over the recording period. Finally, care should be taken to set up the cameras with a proper orientation to ensure optimal overlap of the camera fields of view (Chapter 2). For more detailed information, see our manuscript, "ThruTracker: Open-Source Software for 2-D and 3-D Animal Video Tracking", which will soon be posted to BioRxiv (citation information will be added).

2. Field recording setup and measurements

2.1 Camera selection

ThruTracker is compatible with any camera that saves video to standard formats (.avi, .wmv, etc.). Camera selection will depend on the application. For recording at wind turbines, for example, we have successfully used Flir A65 cameras with 25 mm lenses. These cameras can capture some of the rotor-swept zone as well as detecting animals at heights up to 150-200 m, depending on the size of the animal. We recommend cameras such as the Flir A65 that have dedicated electronic cables for frame-synchronization. Cameras that have audio-in ports can also be used for synchronization, however this option has greater potential for errors. Compatibility for audio synchronization will be added at a later date.

2.2 Camera placement in the field

There is no single ideal arrangement of cameras, as any camera setup represents a trade-off between multiple factors. Closely spaced cameras (Figure 2A) will have fields of view with greater overlap and therefore a larger 3-D volume is captured. However, their views are closer to being parallel, which leads to greater error, especially along the axis of the camera field of view. In contrast, Cameras that are oriented such that their views are closer to perpendicular (Figure 2B) will have narrower regions of overlap between their fields of view and therefore a smaller 3-D volume is captured. However, this orientation will reduce the amount of error in estimating 3-D positions. One rule of thumb that has been proposed is that the baseline distance between cameras should be at least 1/10th of the distance to the

furthest expected target. Therefore, if aiming to capture animals up to 100 meters away, the cameras should be placed at least 10 meters from one another.

2.3 Field measurements

In the field one should measure: 1) the distance between the cameras; and 2) the angle that each camera is inclined relative to the horizon (i.e., inclination or altitude). It is important to take these measurements prior to starting recordings to avoid bumping cameras and changing their orientation. Inexpensive inclinometers with an accuracy of 0.1 degrees are widely available and should be sufficient for most applications.

Figure 2. Comparison of camera configurations optimized for (A) camera overlap and (B) accuracy of 3D reconstruction.

For measuring distances between cameras and between the cameras and the monopole, we recommend using a laser distance measurer such as the Bosch GLM-50C or a similar device that is capable of measuring distances up to 30-50 meters with accuracy within a few millimeters. With practice, it is possible to measure the needed distances with high accuracy during daylight conditions. We recommend mounting the laser distance measurer on a tripod and placing it adjacent to one of the cameras and at the same approximate distance from the other camera. Especially during daylight conditions, you will likely need a reflective target. One can be made by adhering night-reflective tape such as ULINE S-22329W to a 0.3 m x 0.3 m piece of cardboard. A second person should hold the reflective board just above the target camera or on the surface of the monopole. For taking the measurement, it can be helpful to use the edge of the laser measurer as a sighting device. It is recommended to practice this procedure prior to going to the field, starting at shorter distances and progressing to the longest distance that will need to be measured in the field.

For measuring inclination angles, we recommend using a digital inclinometer such as a Wixey WR300 that has a precision of at least 0.1 degrees. Some laser distance measurers such as the Bosch GLM-50C also have a built-in inclinometer that can be used for this purpose. Prior to use, the device should be calibrated by following the instructions in the user manual. One can check the accuracy of the measurements to comparing readings before and after rotating the device 180 degrees on a flat surface. After calibrating and checking the device, one should measure the inclination angle of each camera by placing the device on the flat top surface of the camera in the camera's deployed position.

3. Calibrating camera intrinsic parameters

3.1 Overview

The intrinsic camera calibration determines properties of the camera that allow conversion from a UV pixel on the camera's sensor to a vector originating at the camera and pointing in the direction of the object in the camera's frame of reference. This is an essential component of the 3-D calibration
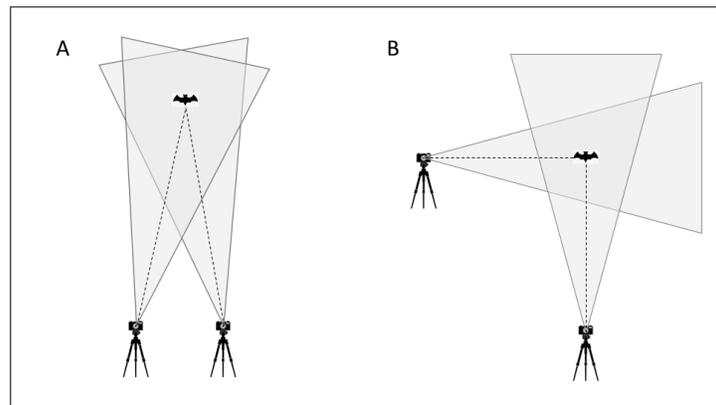
procedure. At a minimum, the intrinsic calibration needs to be done for each camera-lens combination. One may be able to use the intrinsic calibration parameters that someone else generated using the same camera model and lens. However, it is recommended that these values be tested with the individual cameras and lenses by doing a 3-D calibration under controlled conditions. There may be some conditions where an intrinsic calibration is required for each individual camera. For example, when very high accuracy is required or with certain cameras whose production quality is relatively low.

The intrinsic calibration procedure consists of taking multiple images of a checkerboard pattern placed at a known reference distance from the camera. These images should be taken at slightly different 3-D orientations. One complication is that thermal cameras will not detect the visual contrast of a checkboard pattern. Therefore, the pattern must be heated up for thermal contrast to be present. Here, we combine built-in Matlab calibration procedures with a modification of an existing thermal camera calibration procedure (Yahyanejad et al. 2011).

3.2 Setup

Figure 3. Demonstration of intrinsic calibration procedure (left image) and results (right image).
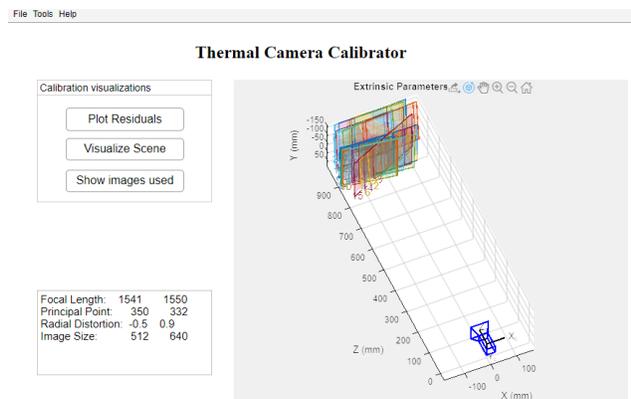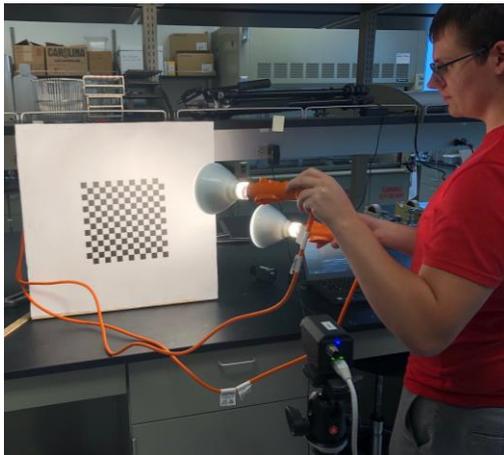
This procedure requires a printed checkboard pattern, camera and recording equipment, a means of heating the checkboard pattern such as two 120-Watt flood lights (white or red lights work equally well), and a meter stick or other device for measuring distances.

3.2.1 The checkerboard pattern
The first step is to determine the minimum distance at which the cameras will provide useable images. The minimum focal distance of the lens can be used as a guideline, but it is best to test this empirically. This can be done by pointing the camera at an object with high heat contrast, such as an A/C adapter that is plugged into the wall. The camera should then be moved toward the object until the minimum distance where the object still provides good contrast. The width and height of the total field of view at this distance should be measured and noted.

Next, a checkboard pattern of the appropriate size should be printed. The total size of the checkboard pattern should be ~80% of the height of the field of view of the camera at the minimum focal distance found in the preceding step. This will allow the full checkerboard pattern to be visible in the field of view at different orientations, and it will be no larger than necessary. It is important to make the checkerboard pattern as small as is acceptable because this will facilitate heating the pattern in a uniform manner (see below). We recommend a 14 x 14 checkerboard pattern (7 black and 7 seven white squares along each axis) with a large white border. The pattern should be printed on poster paper and mounted on a rigid frame. We built a simple frame out of 2 inch by 2 inch wooden boards; another option would be to cut out the center of a poster board. The poster can be held taut and then stapled to the wooden frame with a staple gun. This procedure ensures that the checkboard portion of the poster is not in contact with any other objects. If the poster is mounted directly on another object, the poster will not heat up quickly enough and the images will not be useable.

Figure. 4 Example images of the heated checkerboard pattern (left) and the same image that has been processed with a local adaptive threshold (right).
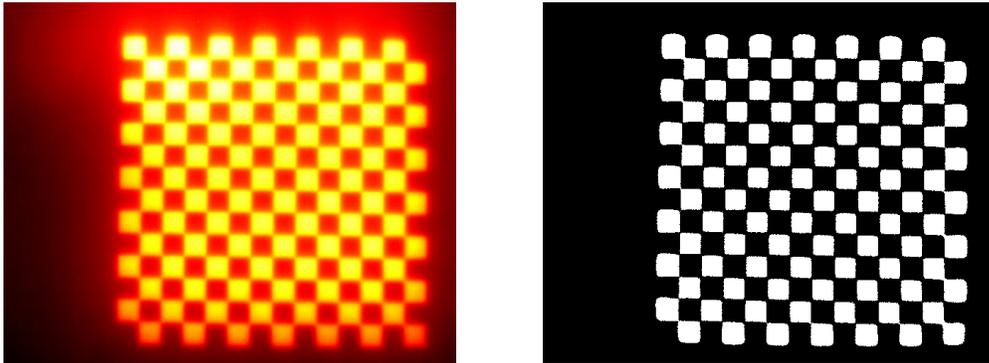
3.2.2 Acquiring images

After printing and mounting the checkboard properly, the camera should be setup so that images can be acquired of the heated checkboard at multiple reference distances. For the Flir A65 cameras, we used reference distances of 0.9 m and 1.0 m with a checkerboard pattern that was a total size of 0.254 x 0.254 m (18.14 mm square sides). One should practice heating the checkerboard pattern with the lights to provide an image that is of relatively uniform illumination/heating. Figure 4 shows an example image that is of sufficient quality for use as a calibration image. We found an effective method of heating the pattern was to move two heat lamps across the image from left to right and then from top to bottom with each pass taking about 2-3 seconds. The heat in the black squares will quickly begin dissipating into the white squares, eliminating the required image contrast. Therefore, the video should be recording the entire time and images should be selected immediately after the pattern has been heated. We found that with good timing, the checkerboard pattern could be moved from left to right across the field of view of the camera such that the pattern could be recorded at different locations within the field of view. We recommend a total of 20-30 images be taken of the pattern at different locations, including some images where the pattern is perpendicular to the camera's field of view, and some images where the image is at an oblique angle up to about 30-45 degrees from perpendicular. After recording several videos of the heated checkerboard at different orientations, the user should find images of acceptable

quality and use the camera's software or other video software to extract .jpeg images of the required 20-30 frames into a designated folder.

3.2.3. Processing the images in thruTracker

A utility within ThruTracker is provided for conducting the intrinsic calibration. One can test this procedure with the test images provided with the installation.

- Open ThruTracker
- Under the Tools Menu select "Intrinsic Calibration". This will open the user interface shown in Figure 3.
- Go to Menu->Tools->Convert Images
- Select the folder where you saved the images in the previous section.
- Converted black and white images (Figure 4) are saved in the subfolder "BW" in the folder where the original images were stored.
- Review converted images and ensure that at least 15-20 underwent a proper conversion. If not, acquire additional images using the procedure in section 3.2.2.
- Under the Menu go to Tools-> Calibrate.
- Enter the number of horizontal and vertical squares and the width/height of each square in millimeters into the dialogue box.
- Review the quality of the calibration by "Plot Residuals" Button (Residuals < 1 pixel indicate acceptable quality).
- Visualize the reconstructed scene by clicking the "Plot Extrinsics" button.
- Examine the scene and check that the re-constructed patterns are at the appropriate distance. If they are either closer or further than they should be, then the focal length will be off proportionally. You may need to add more images at oblique angles to correct for this problem.
- After an acceptable result is obtained, Go to File-> Export Camera Profiles to save the resulting camera parameters file. Note the location of this file for use in generating 3-D Calibrations (Section 4.3).

4. Calibrating camera extrinsic parameters

4.1 Overview

The extrinsic calibration determines the camera positions in 3-D space and their orientations (roll, pitch and yaw). These values, when combined with the intrinsic calibration parameters generated in Chapter 3, allow for calculation of the direct linear transformation (DLT) coefficients. The DLT equation transforms sets of UV points that have been corrected for lens distortion into 3-D points. One can either introduce objects into the scene, such as a wand (Theriault et al. 2014), or use objects that occur naturally within the field of view such as trees, power line poles or animals themselves (i.e., background points, or unpaired points). Because the 3-D positions of background points are not known relative to one another, a relatively large number of points are required. We recommend at least 50-100 unpaired points. The points should be spread out to cover as much as possible of the 2-D view of each camera and the 3-D volume that is captured by the camera set. These points set the overall geometry of the scene.

Figure 5. Example images from two synchronized thermal cameras recording at a wind turbine. A and B show the full field of view from the two cameras. C and D show cropped regions of A and B near the nacelle. Yellow and orange arrows indicate points that were used for the 3-D calibration. E and F show a cropped region around a turbine tip from another point in time.
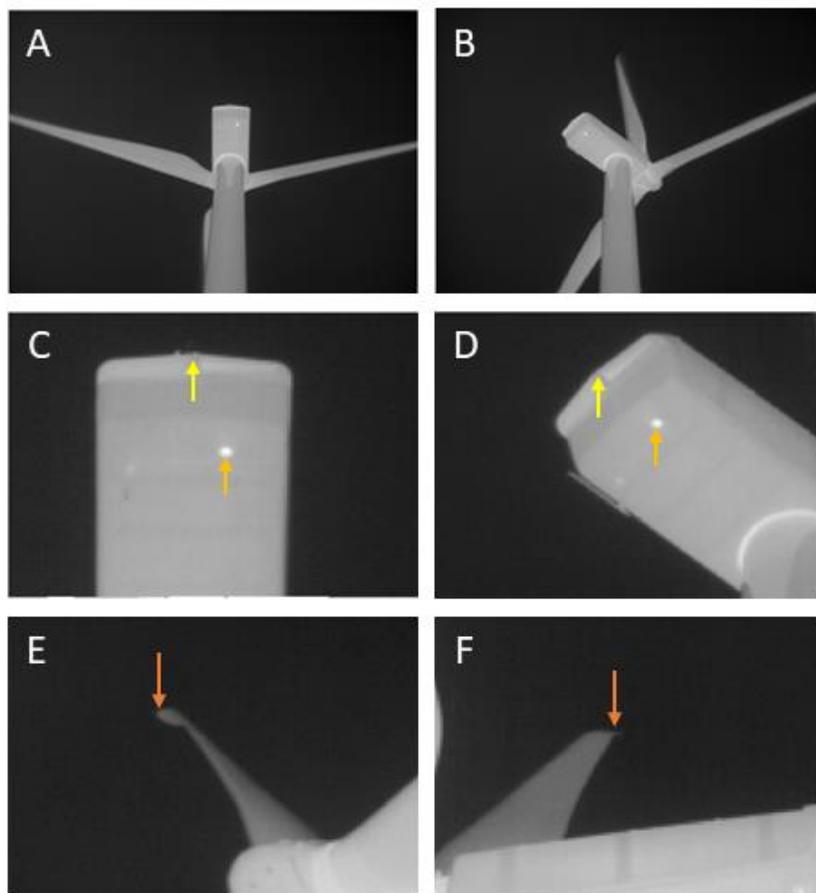
The *scale* of the scene is set by inputting the distance between the cameras, and the scene is oriented to gravity by inputting the inclination of the second (slave) camera relative to the horizon.

4.2 Digitizing background points in DLTdv

Synchronized video should be saved in a standard video format such as AVI or WMV. The videos should be opened in DLTdv8 (freely available for installation on mac, windows and linux machines from http://biomech.web.unc.edu/dltdv/). DLTdv software includes a digitizing platform as well as deep learning modules, which could be used for automatic tracking. More information is available at the link provided above.

In DLTdv8, manually digitize 50-100 points that capture as much as possible of the 2-D space of each camera field of view and the 3-D recording volume. Avoid digitizing the same point or closely spaced points multiple times, as this does not add new data to the calibration and can amplify digitization errors.

One should take great care in finding discrete points that are clearly visible in each camera. For example, Figure 5C and 5D show two points in each of two cameras. The orange arrows indicate a small "hot spot" on the turbine that makes an excellent digitization point. Yellow arrows indicate the pointed center of the back of the nacelle, which also makes a good point. One should avoid using points with curved 3-D edges where it is difficult to match locations between camera perspectives. Figure 5E and 5F show a turbine tip that was used for a 3-D calibration but note that it is difficult to determine the exact matching location of the turbine tip between frames. This is complicated by the curved tip of the turbine blade. One should take care to find images where the 3-D perspective facilitates matching specific points on the turbine and turbine blades. One can also digitize animals flying through the 3-D volume.

After digitizing the points in DLTdv8, click the "Export points" button to save the data for use in simpleWand.

4.3 Making an extrinsic calibration in simpleWand

You are now ready to make a calibration!

- Open simpleWand within ThruTracker by going to the Menu Tools->Make 3D Calibration
- In simpleWand Menu click Project-> New Project
- Enter "2" for the number of cameras
- In the Menu go to Data-> Load Unpaired Points and select the file you exported from DLTdv8
- Click the "Unpaired Points" tab just below the menu bar and check that your points were loaded
- In the Menu go to Data-> Load a camera profile and choose the camera profile provided with the software or that you made using the Intrinsic Calibration procedure in section 3
- Go to the "Alignment and Scale" Tab and enter the last camera's inclination/elevation angle and the distance between the two cameras.
- Go to Menu-> Compute Calibration.
- Review the Calibration scene and check that it is scaled appropriately
- Check the DLT Camera residuals, which are shown just below the Calibration scene. Values < 1 are acceptable, although < 0.5 is preferred.
- Review the Unpaired Point Errors. You may want to remove points that are more than 2-3 times average error rates, but do not remove more than 10-20% of the total points.

- Go to Operations-> Compute Calibration to recompute the calibration.
- When finished save the project (Project-> Save this Project) and export the DLT Coefficients (Project-> Export DLT Coefficients).
- You're done!

Figure 6. Example of a successful calibration using simpeWand. Three clusters of points are visible that correspond to 1) the hub and nacelle (middle cluster), 2) tips of blades pointed away and 3) toward the two cameras (top and bottom clusters, respectively). Note that the hub and nacelle points are at the expected height of ~80 meters above the 2nd (reference) camera indicating proper scaling of the scene. Twelve of 67 points were removed as outliers because they had DLT residuals > 2 pixels. The Mean reprojection error of the remaining points was 0.38 pixels.
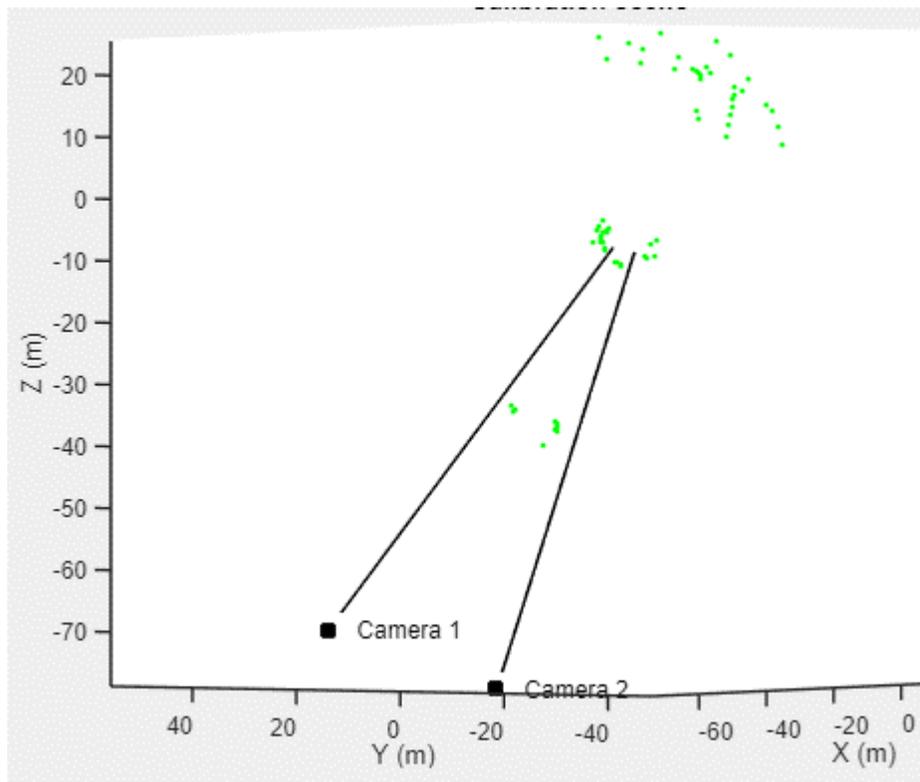
5. Processing video detections

ThruTracker has various settings for 2-D detection and tracking. To process a video go to File-> New 2D Detection. It will open the window shown in Figure 7.

Figure 7. Menu for 2-D tracking options

Detection Options

- Click the "**Choose Video File**" Button and select the video you would like to process. This will automatically populate the "First Frame", "Last Frame" and "Save filename" entry boxes.
- Choose the following settings:
  - **Detector**: Choose either OpenCV or Matlab for background subtraction. We recommend OpenCV, however this option may not be available on certain operating systems.
  - **First and last frame**: Choose the range of frames for processing. You may want to start by processing ~100 frames with known detections to optimize settings before processing longer videos.
  - **Minimum and maximum object pixels**. This can be used to filter detected objects that are outside of this range.
  - **Sensitivity**: Higher values will result in fewer detections, lower values more detections. The desired value will be highly dependent on the particular recording setup. For

OpenCV, recommended sensitivity settings range from 10 to 45. For matlab detector, Sensitivity settings range from 0.6 to 0.99.

- o **Background Frames:** (OpenCV detector only). This is the number of frames used for generating the rolling background image. We recommend a smaller value (~10-50) for rapidly changing scenes and fast moving animals. A higher value (100-500 frames) is recommended for more slowly changing backgrounds and slower animals.
- o **Image Smoothing:** This option applies a [gaussian smoothing filter](#) to the image before detection with a diameter equal to the value entered in pixels. It then [dilates](#) (connects) detected objects that are within half of the image smoothing value. These two procedures enhances detection of objects of a specified size. You should choose a value equal to the width of the smallest object you want to detect. For very small objects, a value of 1 is recommended. You may need to use a software program that allows you to zoom into video images with high magnification to determine the appropriate value. If multiple, close-by animals are being lumped into a single detection, choose a smaller value. If parts of a single animal are being detected multiple times in a single frame, choose a larger value.

- o **Header rows** for your video. Some cameras, such as Flir and Axis include header information on the images themselves. You can use these options to ignore these headers.
- o **Processors.** Choose the number of processors for parallel processing or 1 for not using parallel processing. You can check how many processors are available using the "Check" button. It is recommended to use 1 processor unless you are processing very long videos.
- o **Save filename and location** can be changed if desired. The default is to place the detections file in the same folder as the video file being processed.

2D Tracking Options

After detecting the objects in each frame, ThruTracker attempts to link detections across frames into tracks. These options typically involve a tradeoff between being more conservative in linking points, leading to shorter tracks, but fewer false positives, versus being more aggressive in linking points into tracks.

- **Minimum Track Length**: The minimum number of points that are required to be accepted as a track. Tracks with fewer points than this number will be discarded.
- **Maximum Gap Length:** Maximum number of frames after a detection that ThruTracker will keep looking for another point.
- **Match Threshold**: The number of pixels between the predicted location of the track in the frame and actual position of detections that will be allowed for the detection to be accepted as part of the track.

After setting your options, click on "**Run ThruTracker".**

Reviewing detections

After processing the detections, the data will automatically be saved in the specified file location. It will also be loaded into ThruTracker for visualization and review.

Figure 8. ThruTracker 2-D visualization
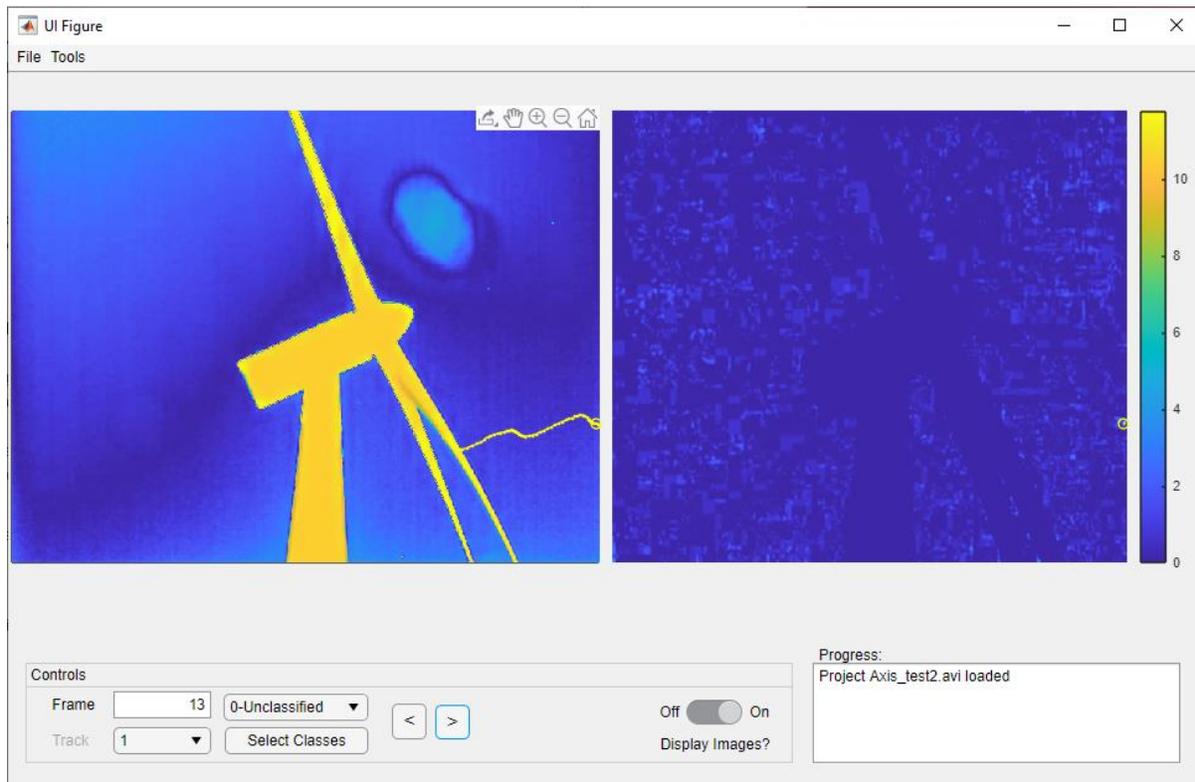
In the **Controls Panel** you can **enter the frame** to visualize, or **advance or go back a frame** using the ">" and "<" buttons. You can also use the left and right arrows on your keyboard to move between frames.

The left image shows a colorized version of the raw image. The right panel shows a simple background-subtracted image. **2-D detections are shown with red circles**. **Tracks are shown as dotted lines**. The current track is shown in red. Other tracks that are present in the current frame are shown in white.

You **can scroll through tracks** either using **the Tracks dropdown**, or by using the **up and down arrows on your keyboard**.

You can export a summary of track data and, or XY detections by Clicking File-> Export Data

Processing Detections for Animal Counts

There are two tools for counting animals. The first is **Track Selector** (Tools-> Track Selector), which allows you to visualize all tracks simultaneuously and categorize them by their positions.

- Either process or load a 2-D detection file.
- Open Track Selector (File -> Track Selector)

- A new window should open with all tracks displayed. The start of the track is indicated with a larger circle.
- Choose from Track Selection dropdown whether to select any portion of tracks or their starts or endings.
- Choose from the "Classify as" dropdown the category for classification.
- Click the "Select Tracks" button
- Draw a rectangle around the tracks you'd like to classify.
- Repeat the above steps until finished with classification.
- Go to File-> Export data to save the classifcations to a spreadsheet.

The second tool for counting animals is **Track Counter.** This tool counts animals based on tracks passing into or out of sections of a rectangle that you draw over the image.

- Either process a video or load a detection file in ThruTracker
- Open Track Counter (File -> Track Counter)
- Draw a rectangle over the image in ThruTracker.
- A pop-up window will allow you to select which parts of the rectangle to include in counts.
- Click "Count!"
- A summary of exits and re-entries will be displayed. Choose whether to save counts to a spreadsheet.

Exporting Videos

There are two options for exporting videos. The first option is to make a cropped **video of individual tracks**:

- Process a video or load a ThruTracker detections file
- Scroll through the tracks until the track you want to make a video of is selected
- Go to Tools-> Make Track Video
- Choose the size of the cropped video in pixels, the video format and frame range.
- The cropped video will be saved to the folder "trackVideos", which will be made in the same folder of the original video that was processed.
- Review the cropped video and re-process if necessary

The second option allows you to make a video that **visualizes all the detections**:

- Process a video or load a ThruTracker detections file.
- Got to Tools -> Make Detections Video
- Choose the frame range
- The detections video will be saved in the same folder as the original video.

6. Generating 3-D tracks

To generate 3-D tracks, you will first need to have made a 3-D calibration for your scene (Chapter 4) and made 2-D detections for each of your video files (Chapter 5). When you are ready to make your 3-D tracks go to the menu File-> Make 3-D Tracks. This will open the dialogue box shown in figure 9.

Figure 9. Dialogue box showing options for generating 3-D tracks


Steps for making 3-D Tracks

- Click the "Select Files" Button
- Choose the number of cameras
- Choose the **detections file** for your cameras
- Choose the **simpleWand or easyWand calibration file**
- Choose the filename and directory if different from the default
- **3-D tracking options:**
  - Choose the frames you'd like to process if different from the range that is auto-populated
  - **Minimum Track Length:** Minimum number of frames required for a track to be accepted
  - **Maximum Gap Length**: Number of frames to keep searching for a detection after the end of a track
  - **DLT Threshold:** Maximum DLT residual error accepted for two detections to be made into a 3-D point
  - **Match Threshold**: Maximum value of combined XY and XYZ errors for a point to be accepted as part of a track. The "Match" value equals the average XY error in pixels for each camera divided by the **XY divisor** plus the XYZ error in meters divided by the **XYZ divisor**.

- Click the "Make 3D Tracks!" button
- A visualization of the 3-D tracks will be shown and the data will automatically be saved in the location specified.

7. Data visualization

ThruTracker includes a built-in data visualization tool called BatVis3 specifically for visualizing tracks at wind turbines. To open the visualization tool, go to the menu Tools-> Visualize 3-D tracks. This will open the BatVis display shown in Figure 10.

Figure 10. BatVis3 track visualization software
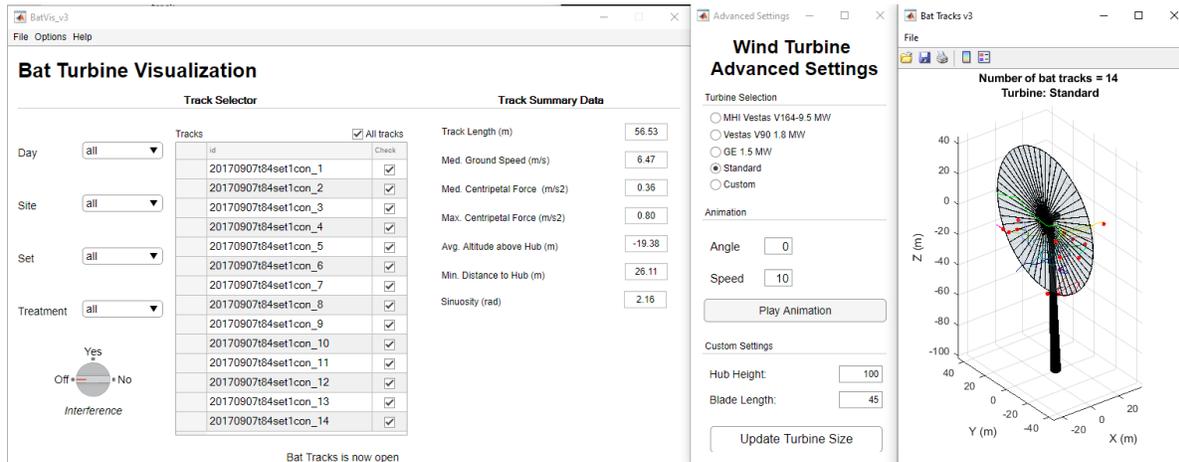
- Load data using the menu File-> Load Data
- Choose the folder containing the XYZ track data and the digitized hub-nacelle points for each track.
- XYZ track data should first be exported from the 3D track file using the File-> Export Data menu in ThruTracker
- Hub-Nacelle points should be digitized in DLTdv8. Each 3-D track should have a corresponding track in DLTdv8. The hub should be digitized using the first frame of each track and the nacelle should be digitized using the second frame for each corresponding track. Note: This workflow will be updated in future versions to thruTracker to speed up the workflow.
- The 3-D track data should be saved using the following format: "bat_date_site_set_treatment_xyzpts.csv".
- The digitized turbine data should be saved using the same format except the file should start with "turbine" instead of "bat".
- After loading the data, the display shown in Figure 10 will be loaded with all of the tracks. The "Bat Turbine Visualization" window allows you to select which tracks are displayed by choosing the Day, Site Set and Treatment
- You can also choose to display only those tracks that fly through the rotor-swept zone (interference="Yes"), or not (Interference= "No") or all tracks (Interference "Off").
- Track Summary Data shows averaged track parameters for all tracks that are selected

Advanced Settings

The Wind Turbine Advanced Settings window allows you to:

- Change the size of the turbine by selecting the turbine model or entering custom dimensions
- Play an animation of the 3-D tracks and animation with options for adjusting the speed of playback and the view angle.
- A view angle of 0 indicates profile view; 45 a 3-D view from partially overhead, and 90 a complete overhead view.

Turbine Visualization

The wind turbine visualization shows the bat flight tracks with respect to the turbine itself (Right panel in Figure 10). The default cursor shows a cross hairs that allows you to select a track. Selecting a track will show track parameters for that individual track. You can also rotate, pan and zoom in and out using the menus at the top of the turbine visualization. Deselect these options to return to the cross hairs for selecting individual tracks.

8. Additional Information and Contact

ThruTracker was developed by Dr. Aaron Corcoran (UC Colorado Springs) and Dr. Tyson Hedrick (University of North Carolina) in collaboration with Bat Conservation International and with funding by the Department of Energy and support from the National Renewable Energy Laboratory.

It is published with a [3-Clause BSD license](.).

We thank William Valentine for contributing to the development of multiple aspects of this software, Michael Schirmacher for guidance and consultation on the development of the software, and Bethany Straw for assistance with obtaining thermal video used for development and testing.

Updates to this software will be available at the following sites:

[www.sonarjamming.com/thrutracker](.)

Questions and comments should be sent to Dr. Aaron Corcoran ([acorcora@uccs.edu](.))

ThruTracker is provided free of charge. Underlying Matlab code will be made publicly available after publication of the software in the peer-reviewed literature. This software comes with no warrantee is meant for use by individuals. It includes components that are published with licenses that prevent its use into additional products that are commercialized.