

Commissioning An All-Sky Infrared Camera Array for Detection Of Airborne Objects

Laura Domine^{1,3,*}, Ankit Biswas³, Richard Cloete^{1,3}, Alex Delacroix^{1,3}, Andriy Fedorenko³, Lucas Jacaruso³, Ezra Kelderman³, Eric Keto^{1,3}, Sarah Little^{2,3,4}, Abraham Loeb^{1,3,†}, Eric Masson³, Mike Prior³, Forrest Schultz^{3,5}, Matthew Szenher³, Wes Watters^{2,3}, Abby White^{1,3}

¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA, USA 02138

² Whittin Observatory, Dept. of Physics & Astronomy, Wellesley College, 106 Central St. Wellesley, MA, USA 02481

³ Galileo Project, 60 Garden Street, Cambridge, MA, USA 02138

⁴ Scientific Coalition for UAP Studies, Fort Myers, FL, USA 33913

⁵ Atlas Lens Co., Glendale, CA, USA 91201

† Head of the Galileo Project

* Correspondence: laura.domine@cfa.harvard.edu

Abstract: To date there is little publicly available scientific data on Unidentified Aerial Phenomena (UAP) whose properties and kinematics purportedly reside outside the performance envelope of known phenomena. To address this deficiency, the Galileo Project is designing, building, and commissioning a multi-modal, multi-spectral ground-based observatory to continuously monitor the sky and collect data for UAP studies via a rigorous long-term aerial census of all aerial phenomena, including natural and human-made. One of the key instruments is an all-sky infrared camera array using eight uncooled long-wave infrared FLIR Boson 640 cameras. In addition to performing intrinsic and thermal calibrations, we implement a novel extrinsic calibration method using airplane positions from Automatic Dependent Surveillance–Broadcast (ADS-B) data that we collect synchronously on site. Using a You Only Look Once (YOLO) machine learning model for object detection and Simple Online and Realtime Tracking (SORT) algorithm for trajectory reconstruction, we establish a first baseline for the performance of the system over five months of field operation. Using an automatically generated real-world dataset derived from ADS-B data, a dataset of synthetic 3-D trajectories, and a hand-labelled real-world dataset, we find an acceptance rate (fraction of airplanes passing in the effective field of view of at least one camera that are recorded) of 41% for ADS-B equipped aircraft, and a mean frame-by-frame aircraft detection efficiency (fraction of recorded airplanes in individual frames which are successfully detected) of 36%. Detection efficiency is heavily dependent on weather conditions, range and aircraft size. Approximately ~500,000 trajectories of aerial objects are reconstructed from this five month commissioning period. These trajectories are analyzed with a toy outlier search focused on large sinuosity of apparent 2-D reconstructed object trajectories. About 16% of trajectories are flagged as outliers and manually examined in the IR images. From these ~80,000 outliers, 144 trajectories remain ambiguous, which are likely mundane objects but cannot be further elucidated at this stage of development without information about distance and kinematics or other sensor modalities. We demonstrate the application of a likelihood-based statistical test to evaluate the significance of this toy outlier analysis. Our observed count of ambiguous outliers combined with systematic uncertainties yields an upper limit of 18,271 outliers count for the five-month interval at a 95% confidence level. This method is applicable to all of our future outlier searches.



Citation: Domine L., Biswas A., Cloete R., Delacroix A., Fedorenko A., Jacaruso L., Kelderman E., Keto E., Little S., Loeb A., Masson E., Prior M., Schultz F., Szenher M., Watters W., White A. Commissioning An All-Sky Infrared Camera Array for Detection Of Airborne Objects. *Preprints* 2024, 1, 0. <https://doi.org/>



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: Aerial anomaly; aerial object tracking; UAP; unidentified aerial phenomena; unidentified aerospace phenomena; unidentified anomalous phenomena; infrared sensors; environmental monitoring; instrument calibration; instrument testing; instrument commissioning.

1. Introduction

There is currently very little publicly available scientific-grade data on Unidentified Aerial Phenomena (UAP) whose properties and kinematics purportedly reside outside the

performance envelope of known phenomena. Several U.S. federal agencies have weighed in on this issue. The Office of the Director of National Intelligence (ODNI)'s annual UAP reports [1,2] have advocated since 2021 for bolstered data collection initiatives and increased resource allocation in order to scrutinize unexplained aerial phenomena that unambiguously pose aviation safety issues and potentially challenge U.S. national security. In 2023, NASA released an independent study[3] which emphasized how studying the phenomenon using passive sensing "with multiple, well-calibrated sensors is paramount" and that "multispectral data" needs to be collected "as part of a rigorous data acquisition campaign". In addition, they suggested that "machine learning algorithms could be incorporated to detect and analyze UAP in real-time." The Galileo Project is designing, building, and commissioning a passive, multi-modal, multi-spectral ground-based observatory to continuously monitor the sky and conduct an exhaustive observational long-term survey in search of measurable anomalous phenomena[4–8]. This long-term field observation effort fits many recommendations from the NASA study and represents "a complex undertaking whose outcome could allow for substantial and systematic gathering of UAP data as well as a robust characterization of the background" [3].

We have designed and built the first observatory at our development site in Massachusetts [4], and we are now commissioning instruments for each of the sensor modalities separately (optical sensors in the infrared, visible, and ultraviolet; acoustic; radio spectrum; magnetic field strength; charged particle count; and weather), and in combination, to validate and benchmark their performance characteristics. Once commissioned, we will commence collecting scientific-grade data and begin to quantitatively identify classes of objects and statistical outliers that are corroborated over time. We can generate testable hypotheses to account for any novel class discovered in this way, which, after further investigation and possible instrument refinements, may result in the discovery of scientific anomalies: i.e., statistical outliers that resist explanation in terms of prevailing scientific beliefs, representing objects and phenomena that are unknown to science [4]. We can also derive upper bounds for the occurrence rate of outliers, novel classes, and anomalies at our development site. As the instruments of our observatory are commissioned, we will make copies and distribute them to additional sites. Such a network of observatories will enable us to monitor a larger volume of the sky. This will not only increase our likelihood of novel class discovery, but also open the possibility of scientific measurement of UAP event rate and spatio-temporal distribution, which to date has never been achieved over significant time and spatial scales.

1.1. Related work

For a review of relevant historical field work and related studies, we refer the reader to section 2.3 of our original design paper [4], and to [9]. A variety of related contemporary work exists, but academic, instrumented, field efforts focused on UAP are rare. These include: the automated, roof-top, and continuously recording system of optical and infrared cameras developed by the Interdisciplinary Research Centre for Extraterrestrial Studies at Julius–Maximilians–Universität Würzburg in Germany [10–12]; a one week field study on and near Catalina Island in 2021, conducted by UAPx, led by a group of academic scientists from the University of Albany, using optical, thermal infrared cameras and high-energy particle detectors [13]; a long-term effort re-purposing data from existing aerial monitoring networks that is underway by "Sigma2" (the UAP technical commission of the learned society, Association Aeronautique et Astronautique de France), which makes use of active and passive bistatic radars, visible and thermal infrared cameras, optical and radio meteoroid surveillance networks (including FRIPON [14], operated by the Institute of Celestial Mechanics and Ephemeris Calculation of the Paris Observatory), gravimeter and magnetometer networks, as well as IR and optical imaging satellites [15]; and the search for fast unexplained transients in photographic plates in historic sky surveys [16], also known as the Vanishing and Appearing Sources during a Century of Observations (VASCO) project, based at Stockholm University.

1.2. Galileo Project commissioning approach

Our observational effort is unique, with its high concentration of sensor modalities designed for long-term operation in outdoor field conditions. We require careful scientific calibration and commissioning to ensure a clear understanding of the reliability, operational environment, and detection volumes of individual instruments and combinations of instruments. Each instrument of our observatory goes through four distinct phases: development, integration and test, commissioning, and science. Commissioning is a final group of activities (generally performed in situ after deployment) that aims to bring a system into an operational condition by ensuring it meets key requirements and science expectations. Meeting basic system requirements is *verification* while confirming science gathering expectations is *validating* the system. Commissioning activities include calibrations, performance characterizations, functionality checks, and data processing evaluations, and are undertaken for both instruments and infrastructure generally on a system level basis, but may occur on subsystems or individual components as needed.

Our approach is primarily driven by a “top-down” paradigm whereby our observatory goals (e.g. an all-sky survey, object triangulation, autonomous data processing with accurate object detection, remote autonomous operations, etc.) taken from our Science Traceability Matrix (STM) [4] will drive observation plans and then specific instrument and operational use cases. These will in turn form the basis of test cases and data collection campaigns that, when evaluated as successful, act as “proof” of performance. Additionally, observatory commissioning will be step-wise on an instrument basis (following the approach commonly used for NASA science missions) with the advantage of collecting scientific data as early as possible. Commissioning is also “end-to-end” from sensing to data product production, meaning that it encompasses the evaluation not only of sensor calibrations but also of final data processing pipelines and the accuracy and utility of the final data products.

Importantly, although we define the completion of commissioning as the demarcation between instrument development and testing vs. scientific data gathering operations, it is typically the case for scientific observatories that commissioning is just the beginning of a continuous improvement process as more is learned about how the instruments and systems perform under in-situ operating conditions. Continued fine-tuning of instruments, improved data handling and processing, better algorithms for science data analysis, new and better data reduction products, better tools for performing remote observatory monitoring and control, etc. are all expected well after completion of initial planned commissioning.

Our all-sky infrared camera array with eight uncooled long-wave infrared FLIR Boson 640 cameras, nicknamed “Dalek”, is the first instrument of our first observatory to enter a commissioning phase. This paper describes the commissioning process and results. In Section 2 we describe the Dalek design, the camera calibrations, and the algorithms to detect and reconstruct object trajectories. In Section 3 we report on commissioning studies of the Dalek and the full system including the detection pipeline. We also propose that UAP instrument-driven studies should adopt likelihood-based statistical tests to evaluate the significance of outlier searches. Section 3.4.2 showcases a *toy* outlier analysis, i.e. a deliberately simplistic outlier search which can be used to illustrate the application of likelihood-based statistical tests to future outlier searches. This simplified outlier analysis searches for reconstructed object trajectories with large sinusity in our commissioning dataset.

2. Materials and Methods

The *reconstruction* of an aerial object transiting in range of our observatory is defined as the process of identifying and characterizing the original object, its properties and kinematics, from the different sensors’ signals and the multiple outputs of the signals analysis pipeline, which can include for example the reconstructed object’s trajectory. This section aims to present our current methods for *reconstructing* aerial objects using solely the infrared camera array whose details are described in Section 2.1. Section 2.2 presents




Figure 1. Left: Mechanical design drawing of the Dalek IR camera array. Right: photograph of the Dalek as constructed at the development site.

the different calibrations applied to the cameras, including intrinsic, extrinsic and thermal calibration. Section 2.3 summarizes the reconstruction of aerial events, for which we use a mixture of machine learning and traditional object tracking algorithms which are evaluated in this section. Also in Section 2.3, we have summarized the datasets referenced throughout the paper for performance evaluations.

2.1. The Dalek infrared camera array

One of the core modalities of the ground-based observatories is a suite of weather-proofed optical sensors which can be used to detect, track, and classify objects in the entire sky in multiple wavelengths: visible, infrared (IR) and ultraviolet (UV). One of the flagship instruments is a custom all-sky array of eight, long-wavelength IR (LWIR) cameras. It is a hemispherical array of seven IR cameras with one additional IR camera pointing towards the zenith, which provide a 360° azimuth by $+80^\circ$ elevation view of the environment as shown in Figure 1. This figure also shows a near-infrared camera ZWO ASI462MC which is built into the Dalek but is not yet being used and thus, is outside the scope of this paper. The prototype has been deployed to our development site in Massachusetts. This site is surrounded by forest and within 5 miles of a regional airport, which ensures a regular flow of airplanes above the Dalek that we can harness for both calibration and commissioning.

The seven IR cameras in the hemispherical array are FLIR LWIR Boson 640×512 (spectral band $7.5\text{-}13.5\ \mu\text{m}$, focal length 8.7 mm, inverse relative aperture of F/1). Their field of view (FOV) is $50^\circ \times 40^\circ$ (horizontal and vertical). The center of their optical axis is pointing 30° above the horizon such that the bottom of the image taken by these cameras corresponds to 10° above horizon. As a result, adjacent cameras have overlapping FOV shown in the left of Figure 2. The zenith camera is a FLIR LWIR Boson 640×512 (spectral band $7.5\text{-}13.5\ \mu\text{m}$, focal length 4.9 mm, inverse relative aperture of F/1.1) with a FOV of $95^\circ \times 72^\circ$. Figure 3 shows a mosaic of the eight cameras views in the array, oriented with respect to true north in map view. Each camera has trees and vegetation in the lower part of its FOV: the same figure also overlays in semi-transparent purple the treeline masks that we use in post-processing to only analyze the sky area of the images. These masks are initially generated using Meta’s recent pre-trained Segment Anything Model (SAM) [17] machine learning model, which classifies every pixel of an image into different semantic categories. The enclosure is a fiberglass dome which provides weatherization. The IR cameras are set behind germanium windows and can be protected from direct sunlight by programmable sunshades. Each camera is connected via a USB3 cable to a Nvidia Jetson Orin NX, which can handle edge computing tasks such as continuous recordings and, in the future, real-time object detection, tracking, and classification. One Jetson receives data from four FLIR cameras. A *Gstreamer* [18] pipeline feeds from each camera and creates 5 minute

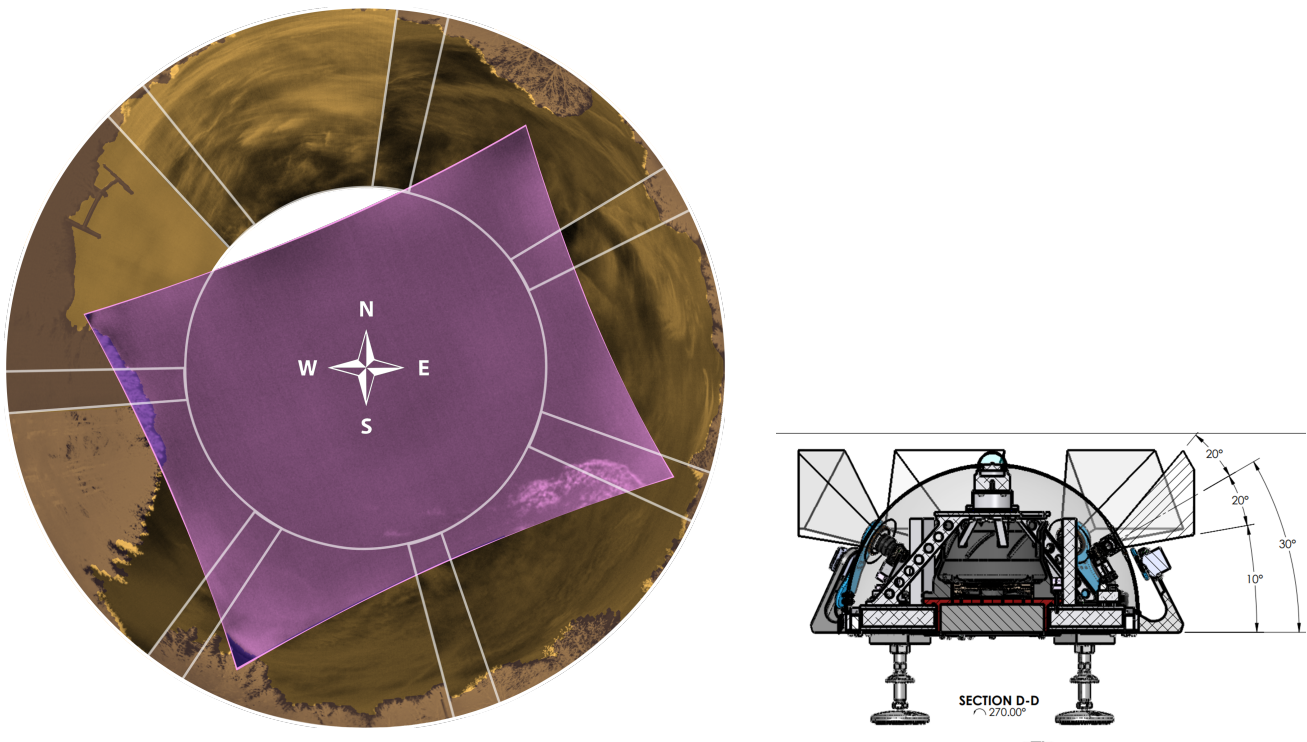


Figure 2. Left: illustration of fields of view (FOV) and their overlap between the eight cameras of the Dalek. The orange areas represent the FOV of the seven cameras arranged hemispherically. The purple area shows the FOV of the zenith camera on top of the Dalek. Right: side-view of fields of view for the hemispherical Dalek cameras. As the center of their optical axis is pointing 30° above the horizon, the bottom of the image that they capture corresponds to 10° above horizon.

video segments which are stored on the Jetson. For all of the data used in this paper, we were recording at 10 frames per seconds, although we plan to increase this to 30 frames per second in the future to take full advantage of these Boson cameras' capabilities. The bit-depth is 8-bit unless otherwise specified. The cameras have a low, high, and automatic gain setting. All the cameras are set on high gain mode. An automatic flat-field calibration (FFC) happens regularly and interrupts the recordings, overlaying a green square in the top right of the frame while the FFC is performed.

The custom Dalek enclosure protects the cameras from weather conditions such as wind, rain, snow, dust, or disturbance from animals. The fiberglass dome has a maximum deformation of less than 1 mm of radial deformations when submitted to a 200 mph wind, according to a finite element analysis of its structure, which is roughly double the record wind gusts in this area. The enclosure contains two, 72 cubic inch, desiccant containers meant to keep the enclosure dry. The desiccant changes color from blue to pink when it reaches its full moisture-absorbing capacity. Baking it releases the moisture and restores the absorbing capacity for a new cycle of use. Finally, a fan in the dome enclosure helps to evacuate hot air from both direct sunlight and hardware heat emission. This keeps the hardware below the maximum operating temperature of 80°C for the FLIR Boson 640 cameras. A first set of germanium windows was chosen for their high transmissivity (95% Anti-reflections / Anti-reflections (AR/AR) coating) over corrosion resistance (85% Diamond-Like Carbon / Anti-reflections (DLC/AR) coating). Prolonged outdoor exposure significantly degraded the AR/AR coating within a year. The germanium windows have since been replaced with ones having both a hydrophobic coating and a high transmissivity ($\sim 96\%$ Hydrophobic / Anti-reflections (HP/AR) coating). These have not shown as many signs of degradation.




Figure 3. Map view of a mosaic of images from the seven hemispheric and the one zenith Boson IR camera, and their orientation with respect to a visible all-sky camera photograph from the Dalek location (background, center of the image). The shaded (purple) semi-transparent overlays show the corresponding treeline masks that are used in post-processing to ignore all but the sky area of the images. All camera frames are taken from a video recording from May 7th, 2024 except for camera 1 (April 3rd, 2024) and camera 7 (May 10th, 2024).

Ideally, all eight of the Dalek's cameras would record continuously 24 hours a day, seven days a week, all year. However, the IR camera sensors can be damaged if the sun is directly in their field of view for long periods, so the Dalek is equipped with a sunshade on each camera that automatically opens and closes on a per-camera schedule as the sun crosses the sky. The mechanical sunshades are mounted on servos and controlled by an Arduino micro-controller. They also have integrated brushes on the side facing the camera to periodically remove debris from the windows. When the sunshade is closed on a camera, its recording is stopped in order to save data storage space. As a result, all cameras run during the night, but during the day, each one has a different recording schedule, with

the north facing cameras recording for longer durations and the south-facing cameras for shorter. Thus, the total expected recording duration is less than 24 hrs/day for each camera. Camera 2 facing north-northwest has the longest recording duration, and camera 6 facing south-southeast has the shortest. The sunshade schedule and total recording time per camera per season is broken down in Table A1. The sunshades are depicted in blue in the schematic on the left side of Figure 1, and are black under red mounts on the photograph of the Dalek as-built, shown on the right of Figure 1.

Although we started recording continuously in November 2023, the system’s uptime only started to stabilize in January 2024. Figure 4 shows the improvement over time of the Dalek cameras’ uptime, quantified by the recording efficiency per day or ratio between the total recording duration and the expected duration from the recording schedule, for each camera. Cameras 2 and 6 suffered the most downtime, and were eventually diagnosed with USB cable failure. After replacement of the USB cable, their behavior stabilized dramatically. The data points in Figure 4 with an efficiency above 8 are due to manual enabling of recording outside of the automatic schedule, for testing purposes.




Figure 4. Stacked area plot showing the evolution over time of the sum of all camera’s recording efficiencies, defined as recording duration per camera per day divided by expected duration based on recording schedule, which varies per camera. If all cameras were recording according to the schedule all the time, the summed efficiencies should add up to 8. The few data points above 8 are due to manual enabling of the recording, for testing purposes. This timeline goes from November 2023 to May 2024. Some cameras such as camera 6 or 7 show a drastic improvement over time.

2.2. Calibration

To relate object positions in the camera image to object positions in the sky, several camera calibration steps are required. Intrinsic calibration corrects for slight aberrations in the lens itself, both intentional (e.g. barrel distortions) and lens imperfections. Removal of image non-uniformities removal (INU) involves correcting for the impact of microscopic lens aberrations on the recorded image under certain lighting conditions. Extrinsic calibration, applied last, enables a mapping from the sensor’s physical location and object’s location in the 2D field of view to the object’s 3D world location, for example from an image’s pixel to an object’s azimuth and elevation relative to camera location.

2.2.1. Intrinsic calibration

The chessboard calibration method described in [19] and coded in the OpenCV library [20] allows us to perform an intrinsic calibration of each camera independently. It yields intrinsic parameters: focal length, optical image center and distortion coefficients. These are stored in the camera configuration file. This calibration is done in the laboratory. Since the




Figure 5. Metal chessboard used for intrinsic calibration of the FLIR Boson 640 cameras. Left: Boson camera image. Center: dark painted metal cutout chessboard on dark painted metal base plate.

Boson cameras are only sensitive to long-wave infrared wavelengths, we cut a chessboard grid out of a metal sheet. To obtain a good contrast between cold and hot, we cool it in a freezer to -20°C , quickly assemble it with thermal insulators as spacers to a solid background metal sheet that has been heated to 50°C , and hold it by a handle that is mounted on the back of the warm metal sheet to capture a series of calibration images covering the field of view of the camera, as shown in Figure 5.

2.2.2. Removal of image non-uniformities (INUs)

The Boson camera manufacturer, FLIR, defines INUs as “a group of pixels which are prone to varying slightly from their local neighborhood under certain imaging conditions” [21]. Specs of dirt, dried water droplets, or microscopic defects in the coating on the outside of the Boson camera’s factory window, or small defects on the inside of the window, can cause INUs. FLIR recommends two procedures to mitigate INUs, lens gain correction followed by supplemental flat-field correction (SFFC), which should be carried in this order [22]. For lens gain correction, we image two metal plates at different but uniform temperatures. The metal plates are covered in black Krylon ColorMaxx paint to ensure thermal uniformity. The software to compute the lens gain correction is provided by FLIR. The SFFC compensates for irradiance from heat generated inside the camera body itself. We occupy the entire FOV of the camera with a blackbody target (polyurethane foam) to perform the SFFC. Both the metal plates for lens gain correction and the black foam for SFFC, are pictured in Figure 6.

2.2.3. Extrinsic calibration with ADS-B equipped airplanes

The rotation matrix and translation vector which convert a camera’s coordinate frame to world coordinates constitute the extrinsic calibration parameters of a camera. To find the translation vector, we use GPS coordinates from a cellphone positioned next to the Dalek, which has an accuracy of approximately 5 meters [23] in this case. Extrinsic calibration of visible light cameras often relies on astrometry, but the LWIR Boson cameras cannot see stars. Therefore, to find the rotation matrix, we adopt the novel calibration technique described in [4] using airplanes, which emit at LWIR wavelengths and also reflect emissions from the warm ground. Automatic Dependent Surveillance Broadcast (ADS-B) Mode-S transmission devices are operated in most airplanes flying in US airspace, as required by the US Federal Aviation Administration [24]. Airplanes send with ADS-B, in near real-time (transmission of the aircraft position must happen within 2 s from the measurement time, with up to 0.6 s of uncompensated latency, and at least once a second [25]) and roughly once every two seconds, their GPS-derived latitude and longitude positions, and GPS and barometric-derived altitudes. ADS-B receivers are widely available for retail and anyone with a software-defined radio receiver can collect local ADS-B data. Non-profits such as OpenSky Network aggregate ADS-B data, crowdsourcing individuals’ ADS-B receiver data from all over the world into a single database [26]. This historical ADS-B database is made available to academic researchers. For this paper, we rely on the OpenSky historical database [26] for ADS-B records to obtain the coordinates of the many aircrafts flying in




Figure 6. Left: Four cameras mounted on a fixture and pointed at the black metal plate used for removing INUs through lens gain correction and supplemental flat-field correction (SFFC). Middle: Four cameras mounted on a fixture and pointing at the black polyurethane foam in the incubator; used for thermal calibration of four Boson cameras at the same time. Right: example of 16-bit image taken with a Boson camera using the thermal calibration setup described in Section 2.2.5. The two small pieces of reflective tape, pinned at the top and left side of the foam, are used to find the geometrical center of the foam in the image, where the target temperature was measured using a thermometer. The “ears” of foam minimize thermal reflections from the sides and door of the freezer or incubator.

the Dalek’s field of view because our own ADS-B receiver was not operating continuously during the five-month period of the Dalek commissioning.

If the camera captures N images of an aircraft as it passes overhead, we can interpolate the 3D coordinates provided by the ADS-B data over this time range, and associate a 3D world location with the aircraft’s camera pixel location at the timestamp of each frame. Let $W = (w_i)_{1 \leq i \leq N}$ be the vectors pointing from the camera’s known origin in world coordinates, given by GPS, to the airplane’s positions in the same frame as given by ADS-B. We then use OpenCV’s implementation of the Perspective-n-Point (PnP) algorithm [27] to compute the camera’s pose in world coordinates, giving as input the camera’s intrinsic parameters, extrinsic parameters, the airplane image points and corresponding world pointing vectors. This returns a rotation matrix which represents the camera’s pose. This method to extrinsically calibrate a camera can be applied to both optical and IR cameras. Applying it to the Dalek data was our first demonstration of this method, and we plan to use this for other cameras in the observatory.

2.2.4. Monitoring camera orientation changes

We automated this extrinsic calibration technique in order to repeat this at regular intervals. This opens the door to detecting unwanted changes in camera orientation: for example, this can happen during high wind weather events or hardware maintenance that requires opening the dome, or temperature-related expansion of the aluminum bench supporting the cameras. Small changes affect the treeline mask used downstream in the analysis pipeline for reconstruction of aerial events, and can cause incorrect estimation of object positions in the sky, which depend on the calibration matrices. Determining this threshold and using this to detect significant changes is important for consistency and accuracy in tracking and trajectory reconstruction.

The process we use to monitor for unexpected changes in camera orientation is as follows. We use the You Only Look Once (YOLO) machine learning (ML) object detection method [28,29], described in Section 2.3, to automatically detect aircraft in a given frame of a Boson camera video. The existing extrinsic and intrinsic calibration allows us to translate an aircraft’s ADS-B GPS coordinates into camera 2D image coordinates. The ADS-B predicted positions are interpolated in a time interval of ± 5 seconds, to account for timing




Figure 7. Euler angles representing the orientation matrix of camera 1 in the Dalek over a period of three months, January to March 2024. The error bars are estimated from the inverse Hessian matrix after minimization. Large error bars are correlated with a lack of statistics. Overall the fluctuations are of the order of 1%.

inaccuracies in both the ADS-B and the camera image frame timestamps, and matched to the closest center of the detection bounding boxes predicted by YOLO. For matching we use a distance threshold of 20 pixels and a range threshold of 2 km from the observatory to the aircraft. We allow such a generous time threshold to account for the occasionally less-than-ideal timestamp accuracy of the video files recorded during this commissioning phase. The distance threshold accommodates latency-induced inaccuracy in the ADS-B derived position estimation and inaccuracies in the existing extrinsic calibration. We collect data from a minimum of 20 airplanes before computing a new orientation matrix. This procedure is repeated on a weekly basis to monitor the changes over time. Figure 7 shows the actual evolution over time of the rotation matrix of one of the Dalek cameras. The rotation matrix is represented by its three Euler angles and computed over a period of at least 1 week, ensuring statistics from at least 20 airplanes. We observe slight fluctuations of the Euler angles over time, of the order of 1%. Errors are derived from the inverse Hessian matrix, obtained using the minimization algorithm [30]. These errors do not account for the uncertainty of bounding boxes center positions, which could affect this measurement. This method will be improved with further refinements but it can be used in its current form to trigger an alarm if the orientation of a Boson camera changes significantly, for example by more than one degree in any Euler angle.

2.2.5. Thermal calibration

The FLIR Boson 640 cameras have the ability to record images with 8- or 16-bit depth. The 16-bit data can be used for thermal radiometry. The camera model selected for the Dalek comes without factory thermal calibration, hence we performed the thermal calibration of the cameras ourselves. This requires taking images of a black-body target of known emissivity when both the camera and the target are exposed to a wide range of temperatures. This is necessary in order to correlate the 16-bit pixel values and camera temperature to the target temperature.

Our target was a block of black polyurethane foam. In order to calibrate up to four cameras at the same time, we used the custom-built mechanical fixture that holds both the cameras and the germanium windows, shown previously in Figure 6 and used for the intrinsic calibration procedure. Since the opacity of the germanium windows is temperature-dependent, being almost opaque at 100°C, it is important to include them in the calibration fixture. We drilled holes in a plywood sheet to keep the distance and orientation of the fixture to the foam constant. Two pieces of foam were taped into position, one inside an

incubator and one inside a freezer. Figure 6 shows an example of this setup, where the fixture faces the foam target in the incubator. A thermal probe is inserted in the foam to measure the temperature at the center of the block. The position of the probe tip inside the foam was marked using horizontal and vertical reflective tape indicators at the edge of the foam. We used the temperature sensor onboard the camera chips to record the camera temperatures. We explored a 2D grid of target and camera temperatures in the range of $[-20, 60]^{\circ}\text{C}$. For a given target temperature, the process was as follows:

1. Warm or cool down the foam to the target temperature.
2. Warm up the camera to 60°C using an incubator.
3. Place it in front of the foam target and capture images at regular intervals while it is cooling down back to room temperature.
4. After cooling down the camera to -20°C with the help of a freezer, place it in front of the target and capture images while it is warming back up to room temperature.

Ambient humidity proved to be a challenge in this process, as it created condensation droplets on the camera lenses and germanium windows. We had to maintain the air humidity in the lab below 40% in order to take meaningful measurements. Figure 8 shows the temperatures sampled for different cameras following this process. Once all the images are taken, i.e. once sampling the 2D grid of camera and target temperature is completed, we find the pixel coordinates of the centers of both foam targets (the foam in the incubator and the foam in the freezer) for each individual camera. We record the mean pixel value in a 32×32 pixel square at these coordinates. We grid the pixel values using inverse distance weighting interpolation and 3×3 smoothing. The result is fitted to a bi-quadratic polynomial, where the two variables are $x =$ camera sensor temperature and $y =$ target temperature and the output is $z =$ pixel value. Figure 8 has an example of fitted polynomial that can be used to relate a 16-bit pixel value and camera temperature to the target radiance, assuming it is a blackbody. The black-body emission of the target is related to its temperature through Planck's equation, which we numerically integrate for the wavelength band of the Boson camera i.e. from $\lambda_1 = 7.5\mu\text{m}$ to $\lambda_2 = 13.5\mu\text{m}$:

$$\mathcal{L}(T) = \frac{\epsilon}{\pi} \int_{\lambda_1}^{\lambda_2} d\lambda \frac{C_1}{\exp(C_2/(\lambda T)) - 1} \lambda^{-5} \quad [\text{W}/\text{m}^2/\text{sr}]$$

where $C_1 = 3.7415 \times 10^8 \text{ W}\cdot\text{m}^2\cdot\mu\text{m}^4$ and $C_2 = 1.43879 \times 10^4 \mu\text{m}\cdot\text{K}$.

2.2.6. Object temperature measurement

This section demonstrates a first-order example of using the thermal calibration of the previous section to estimate the temperature of an aerial object. A surface emits thermal radiance depending on both the surface temperature and the surface emissivity. To use the calibrations from the previous section on real images captured with a Boson camera, we need first to measure the emissivity $0 \leq \epsilon \leq 1$ of the foam used in the calibration process. We let the foam equilibrate to room temperature. We measure the temperature of the foam through a surface temperature probe which yields $23^{\circ}\text{C} \pm 0.5^{\circ}\text{C}$. Simultaneously, an external laser infrared thermometer set to an emissivity of 1.0 measures $22^{\circ}\text{C} \pm 1.5^{\circ}\text{C}$ at the site of the temperature probe. We deduce an estimate of the foam emissivity from the ratio of the two temperatures in Kelvin units, $\epsilon = 0.997 \pm 0.007$.

For the thermal calibration we assumed that, owing to the short distance between the camera and the foam target, any atmospheric effects would be negligible. In general, however, another parameter is involved in the measurement of an object's temperature through our cameras in the field: the atmospheric transmission, which quantifies the attenuation of thermal radiation emitted by the object due to water vapor and carbon dioxide along the transmission path. This depends on the relative humidity, the temperature of the atmosphere, and the distance from the camera to the object. We use the Passman-Larmore tables [31,32], which were derived experimentally, to compute the atmospheric transmission coefficient $0 \leq \tau \leq 1$. For a wavelength of $\lambda = 10.5\mu\text{m}$ (the center of the FLIR Boson camera spectral band), a distance $d = 5000\text{m}$ to the object, relative humidity

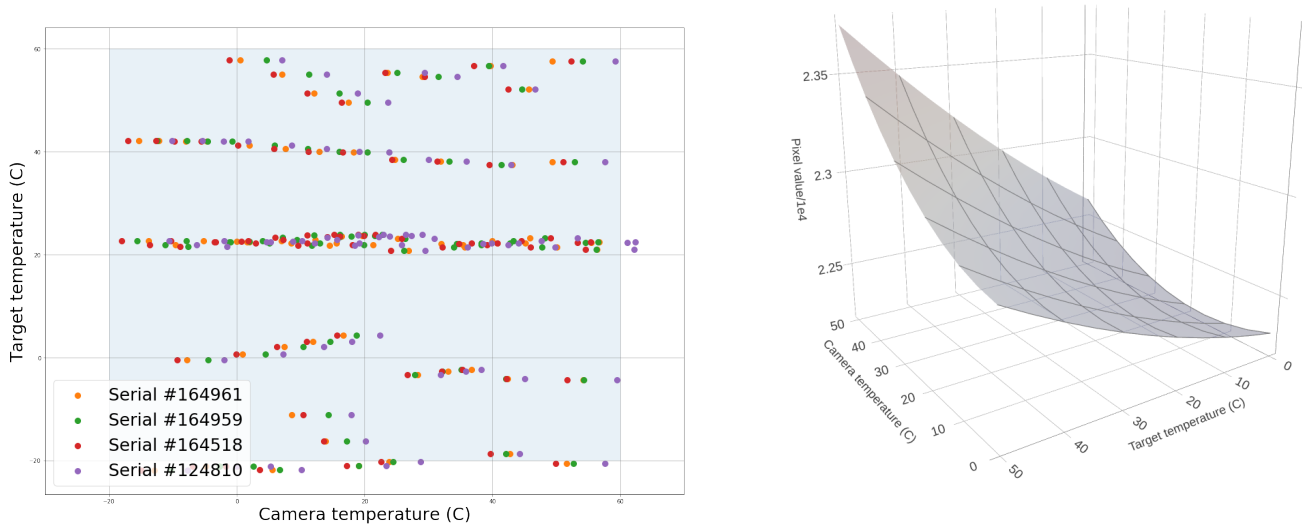


Figure 8. Left: map of measurements made on four cameras of the Dalek. Right: visualization of fitted polynomial. The maximum pixel values occur when the camera and target are both “hot”. Pixel values are divided by 10^4 on the z-axis.

$\omega = 44\%$, and atmospheric temperature $T_{atm} = 22^\circ\text{C}$, we use Equation (6), Table 1 and Table 2 from [32] to compute the water vapour absorbance $P_{H_2O} = 0.544$ and carbon dioxide absorbance $P_{CO_2} = 0.978$. Then we can approximate the atmospheric transmission coefficient with [33] $\tau \sim P_{H_2O}P_{CO_2} \sim 0.532$.

Given the emissivity of the object ϵ and the atmospheric transmission τ , we see a brightness of

$$\mathcal{L}_{\text{observed}} = \epsilon\tau\mathcal{L}_{\text{object}} + \tau(1 - \epsilon)\mathcal{L}_{\text{ambient}} + (1 - \tau)\mathcal{L}_{\text{atmosphere}}$$

where $\mathcal{L}_{\text{ambient}}$ is the brightness emitted by the object’s surroundings and reflected by the object, $\mathcal{L}_{\text{atmosphere}}$ is emitted by the atmosphere, $\mathcal{L}_{\text{observed}}$ is what we measure from the camera sensor, and $\mathcal{L}_{\text{object}}$ is emitted by the object, i.e. the quantity we wish to determine. [32,34] Figure 9 shows an example of 16-bit frame from camera 3 of the Dalek during an airplane fly-by. The aircraft is an A321 taking off from BOS airport at a slant range of 5.3 km from the observatory. We measure the ambient reflected brightness in units of pixel values, and then use our calibration to convert to temperature. We average the pixel values immediately surrounding the object. For measuring the atmosphere brightness, we average the pixel values in the top left quarter of the frame, which is empty sky. We obtain 19748 and 19765 respectively. By plugging in the values for the object emissivity and atmospheric transmittance, we can infer the actual object brightness 22672 from the observed brightness 21307. This would be the pixel value for a piece of foam in the same conditions as it was during the thermal calibration. The fitted relationship between brightness and radiance that we obtained from the calibration in the previous section yields an estimate of the object’s radiance: assuming a sensor temperature of 30°C given the ambient temperature at the time the image was taken, the target radiance is estimated at $53 \text{ W}/\text{m}^2/\text{sr}$. Next we set equal the radiant flux emitted from a surface of $\sim 8 \times 8\text{m}^2$ (assuming an average airliner has a wingspan of 50m on average, at a distance of 5.3km, one pixel of the camera corresponds to $\sim 8\text{m}$ of the object), which sees the camera sensor from a solid angle of $\sim 7.10^{-6}\text{sr}$ at $\sim 5\text{km}$ distance, to the radiant flux emitted by an object at $\sim 0.1\text{m}$ distance, with $\sim 0.25\text{m}^2$ surface area, and which sees the camera sensor from a solid angle of $\sim 0.001 \text{ sr}$: we find the




Figure 9. Left: Example of 16-bit infrared image of an aircraft taken with a Boson camera. Right: zoom onto the airplane in the scene. The red lines cross at the pixel of maximum brightness. The area in between the green rectangles is used to estimate the ambient brightness.

equivalent radiance of the actual target, the aircraft at 5km, to be $\sim 37 \text{ W/m}^2/\text{sr}$. Planck’s law, assuming an emission coefficient of $\epsilon = 0.7$ for an aluminum and stainless steel object, yields a corresponding temperature of 24°C . We compare this result to a temperature estimate based on aircraft altitude. The standard adiabatic lapse rate is a temperature decrease of 2°C for every 1,000 feet increase in elevation, so we expect a temperature of $\sim 2^\circ\text{C}$ at the aircraft barometric altitude of 11,450 feet. Since the aircraft just took off, to estimate its body temperature we use Newton’s law of cooling for an aluminum (specific heat capacity 903 J/kg/K , heat transfer coefficient $\sim 5 \text{ W/m}^2/\text{K}$) sphere of ~ 83 metric tons going from a temperature on the ground of $\sim 20^\circ\text{C}$ to the temperature of 2°C at the aircraft current altitude within 5 minutes. We find that we should expect a temperature of $\sim 18^\circ\text{C}$ for this specific aircraft, which would mean an error of $\sim 6^\circ\text{C}$ in this example.

This temperature estimate is included in this study as a proof-of-concept demonstration. During the commissioning period we collected only 8-bit recordings and did not yet have triangulation in place for range estimation. The rest of this paper is not concerned with objects’ temperature. In the near future, we plan to enable the continuous recording of 16-bit infrared videos with the Dalek, which will help characterize the temperature of objects in combination with range estimation derived from triangulation, for example.

2.3. Reconstruction of aerial objects using YOLOv5 and SORT

The primary design purpose of the Dalek system is object detection and trajectory reconstruction. A combination of machine learning and traditional algorithms suited for this purpose are applied to the IR images obtained from the Dalek video recordings. You Only Look Once (YOLO) [28,29] is a common machine learning architecture for real-time object detection. We train the version 5 of YOLO using both synthetic images and the Dalek’s real-world images [7]. The datasets are described in Section 2.3.1 and the training outcomes are evaluated in Section 2.3.2. YOLOv5 marks detections on each frame individually. To reconstruct the trajectories of detected objects within a single camera’s field of view, we apply a Multi-Object Tracking (MOT) algorithm to track individual objects in a video sequence. We use the popular Simple Online Real-time Tracking (SORT) [35] which combines a Kalman filter with the Hungarian algorithm. Section 2.3.3 contains more details about our implementation and a performance benchmark for this stage.

2.3.1. Datasets for training and evaluation

Synthetic images dataset

An object detection algorithm such as YOLOv5 requires supervised training and tens or hundreds of thousands of labelled training samples. We developed a synthetic image generation tool called *AeroSynth*, which leverages the Python bindings of Blender, a free and open-source 3D modelling software. [7] contains more details about *AeroSynth* and the

training process of a YOLOv5 [28,29] architecture using synthetic datasets. This synthetic dataset contains 25 object bounding boxes (labels) per image, sampled from 40 different 3D models including airplanes, balloons, drones, birds, and helicopters, but excluding some natural objects such as leaves and clouds, totalling $\sim 800\text{k}$ objects spread over $\sim 32\text{k} + 8\text{k}$ (train + test) images.

Mixed synthetic and real-world images dataset

In addition to bootstrapping the training using a synthetic dataset, we also selected real-world images from Boson cameras which contained highly confident (score > 0.9) detections from the YOLOv5 model trained using exclusively synthetic data. We included these images, using these detections as truth bounding boxes or labels, to fine-tune the model on real-world images as well. This mixed dataset contains $\sim 424\text{k}$ objects spread over $\sim 45\text{k}$ real-world images. Roughly 10% of these images are background, empty images. On average, there are 9 object bounding boxes (truth labels) per image. In Section 2.3.2 we compare these two training approaches (using the synthetic dataset and using the mixed dataset) in a benchmark.

Synthetic videos dataset

A video is a set of frames or images that are correlated in time. In order to evaluate tracking algorithms on synthetic, controlled trajectories, we need to generate consecutive images of synthetic objects that, when stitched together, form a video of object trajectories. We again rely on AeroSynth [7] to render videos of 3D models that mimic the video recordings by individual Boson infrared cameras. Each synthetic video is a collection of 100 images or frames, which are snapshots of the 3D models in the scene at successive points in time, separated by 0.1 second. The rate of 10 frames per seconds is chosen to match the Dalek commissioning dataset frame rate. We include the same 3D models that were used to create the synthetic images dataset used to train YOLOv5. The synthetic trajectories fall in three categories: straight, simple curve, and piecewise. Straight trajectories are given a starting point chosen randomly in the 3-D space within the field of view of the camera and follow a straight line in a random, uniformly sampled direction. Curved trajectories are given two points chosen randomly in the 3-D space within the field of view of the camera and follow a circular arc between the two points using a random radius uniformly sampled in the range of $[0, 50]$ meter. Piecewise trajectories are given a starting point in 3-D space within the camera field of view and follow a sequence of straight lines in randomly selected directions. They can have 1 to 5 inflection points. For all trajectories, the speed is uniform across the trajectory and sampled in the range of $[0, 100]$ meters per second, which means that the trajectories points are uniformly sampled. We generate ~ 1600 unique trajectories. Each video contains up to 5 moving objects at the same time. Figure 10 shows visualizations of each type of trajectory.




Figure 10. Comparison of straight (left), curved (middle), and piecewise (right) 3-D trajectories, from the synthetic videos dataset. Trajectories are identified by color and unique number.

Manually labeled real-world images dataset

We manually labeled each frame of a subset of Dalek video recordings by drawing bounding boxes on visually identified aerial objects in each frame (airplanes, birds, leaves,

etc) to create truth labels for object bounding boxes, using the free and open-source Computer Vision Annotation Tool (CVAT). For each of the eight cameras, we sampled up to three five-minute videos per day during January 2024: one at midnight, one at noon, and one at 6 PM Eastern Standard Time. Data availability determined the final dataset, comprising 314 videos, which amount to 904,257 individual frames. A total of 36,036 frames in this dataset (approximately 4%) contain objects. Of these, approximately 94% contain one object, 5% contain two objects, and 1% contain three or more objects. The highest number of objects observed in a single frame is 10. Ultimately, this dataset contains 40,268 individual object annotations (labels) which include the bounding box center coordinates, width and height, but not object class. Throughout the process, we strove for tight bounding box definitions with minimal surplus margin, consistent labeling of all visible objects, and complete enclosure of objects within bounding boxes (even in cases of partial object occlusion).

ADS-B derived real-world dataset

We expect the fraction of aircraft not transmitting ADS-B to be small at our site (military aircraft only infrequently fly overhead). Using ADS-B data (see Section 2.2.3) from the OpenSky database, we select ADS-B records whose latitude and longitude are within a N-S, E-W square of side 10 km centered on the observatory. Additionally, for each selected ADS-B record, we use the camera’s extrinsic and intrinsic calibration to translate the 3D coordinates of the aircraft to 2D coordinates in the frame of the camera images. We can determine that way whether the aircraft’s transmitted ADS-B location is within the field-of-view of at least one camera, and if so, whether it is above the treeline of the camera in which it is visible. The timestamp of the ADS-B record can be compared to the Dalek recording periods to determine whether the camera was recording when the aircraft was in view. We note that in addition to inaccuracies in the camera recording timestamps, there are also latencies in the ADS-B records timing and update frequency that require relaxing any time matching thresholds when working with this dataset: for example, [36] found a mean delay between the aircraft on-board time and the ground receiving time of 0.2 s, which corresponds to 2 frames in the Dalek recordings, while [37] reported that 13% of the ADS-B updates happened at intervals greater than 2 s. This dataset thus contains the trajectories of all aircraft that passed over the site during the five-month commissioning period, and for each of them, which cameras both were recording and had the aircraft in their unobstructed field of view (if not necessarily visible or detectable) at that time. In addition, the hourly weather information corresponding to every aircraft ADS-B transmission is downloaded from the Open-Meteo API [38] and joined to the dataset.

2.3.2. YOLOv5 benchmark on manually labeled dataset

The first stage of the Dalek detection pipeline is a frame-by-frame object detection machine learning algorithm, YOLOv5 [28,29]. Here we present some basic metrics on the performance of the YOLOv5 model that we use for the rest of this paper. In particular, we explain why we decided to use the YOLOv5 model trained on the synthetic dataset alone rather than the model trained on the mixed synthetic and real-world dataset. We apply both of the YOLOv5 models in evaluation mode, i.e. run their inference, on each frame of the 314 videos in the manually labeled dataset of Dalek real-world images. Next, we calculate the *Intersection over Union* (IoU) between each detection bounding box and each ground truth bounding box. The IoU is a ratio of areas which measures the overlap between the predicted bounding box and the ground truth bounding box, providing a metric for the agreement between detections and ground truth bounding boxes. We set a low IoU threshold of 0.01 to define a match between *predicted and ground truth bounding boxes*: i.e. a true positive. We consider detections with an IoU below this value as false positives. Following the IoU calculation, we evaluate the detection performance using the following key metrics:

- *Detection errors*: True Positives (TP), False Positives (FP), True Negatives (TN), and False Negatives (FN), in counts

- *Precision (Prcn)*: the ratio of true positives to the sum of true positives and false positives, $Prcn = \frac{TP}{TP+FP}$
- *Recall (Rcll)*: the ratio of true positives to the sum of true positives and false negatives, $Rcll = \frac{TP}{TP+FN}$
- *Accuracy*: the ratio of correct detections (both true positives and true negatives) to the total number $TP + FP + TN + FN$
- *F1-score*: the harmonic mean (appropriate for finding the average of two rates) of precision and recall

Table 1. Comparison of metrics for YOLOv5 models trained on real-world + synthetic data versus synthetic-only data

Metric	Real-World + Synthetic Data	Synthetic-Only Data
True Positives (TP) count	2,588	24,465
False Positives (FP) count	435	8,728
False Negatives (FN) count	36,412	14,535
Precision	85.6%	73.7%
Recall	6.60%	62.7%
Accuracy	6.60%	51.3%
F1-Score	12.3%	67.8%

As illustrated in Table 1, the model trained on both real-world and synthetic data achieves a high precision of 85.6%, indicating that when it makes a positive prediction, it is likely correct. However, its recall is extremely low at 6.60%, meaning it fails to detect a large portion of the actual objects, as evidenced by the high number of false negatives (36,412). This poor recall significantly impacts the model's overall accuracy (6.60%) and F1-score (12.3%), suggesting that this model struggles with generalization or may be overly conservative in making detections. In contrast, the model trained solely on synthetic data shows a much better balance between precision and recall. While its precision (73.7%) is slightly lower than the mixed-data model, it achieves a much higher recall (62.7%), leading to a significantly improved F1-score (67.8%) and accuracy (51.3%). This model detects a greater number of true positives (24,465) and substantially fewer false negatives (14,535), indicating it is more effective at identifying objects, though it also produces more false positives (8,728).

Given these results, we hypothesize that for larger objects, the model trained on a mix of real-world and synthetic data might perform better than its synthetic-only counterpart. Indeed real-world data could provide more accurate representations of larger objects, which might be underrepresented or inaccurately simulated in synthetic datasets. To further explore this, we repeat the inference multiple times for each model, adjusting the bounding box area minimum threshold incrementally from 0 to 100 pixels square. The results, as shown in Figure 11, reveal how the performance metrics—accuracy, F1-score, precision, and recall—vary as the minimum bounding box area increases.

For the model trained on both real-world and synthetic data, accuracy and recall improve significantly as the minimum bounding box area increases, with the most notable gains occurring between thresholds of 40 and 80 pixel square. This suggests that the model becomes more reliable when detecting larger objects, which aligns with our hypothesis. The F1-score also increases steadily, reflecting the improved balance between precision and recall as the model becomes more selective in its detections. On the other hand, the model trained solely on synthetic data shows relatively stable performance across varying bounding box thresholds. While precision remains consistently high, recall improves slightly with higher thresholds, resulting in a stable F1-score. The accuracy of this model, however, begins at a higher value and increases more gradually, indicating that it is generally more consistent in detecting objects of varying sizes, but it might not show as much performance increase from increased object size as the mixed-data model does.

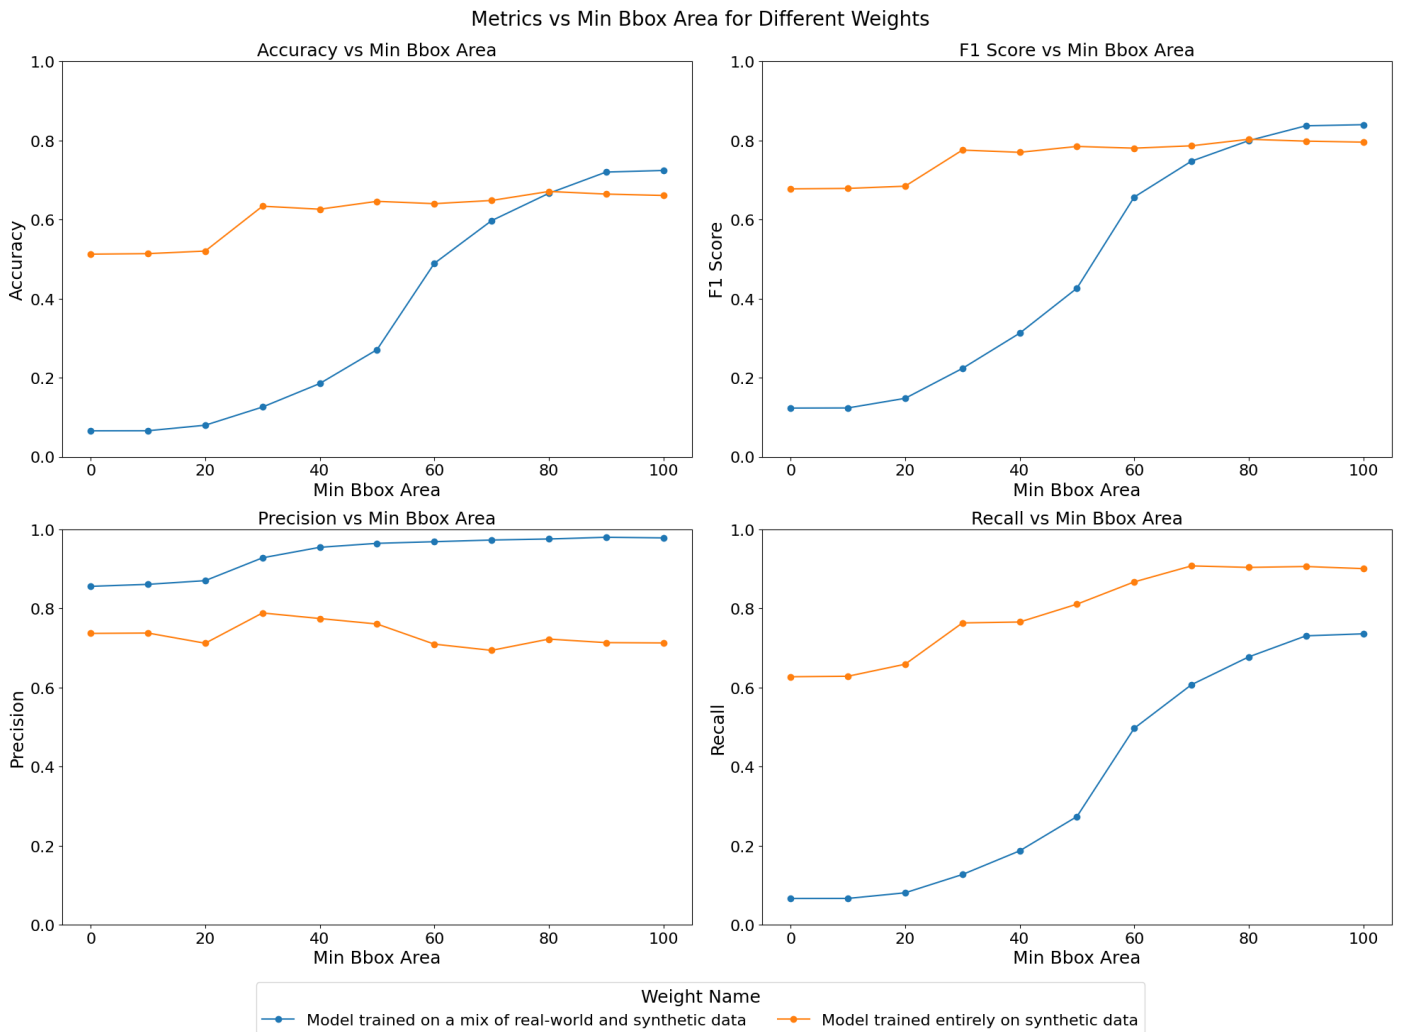


Figure 11. Comparison of YOLOv5 model trained on synthetic-only dataset versus model trained on a mix of synthetic and real-world datasets. The horizontal axis units are pixel square. The model trained on synthetic-only data has better overall performance and is the one used for commissioning in this paper.

Overall, the synthetic-only model demonstrates superior performance across most metrics, particularly in terms of recall and overall accuracy, making it a more reliable choice for object detection in this context. All further studies involving YOLOv5 in this paper use the synthetic-only model. In the future, further tuning and possibly the inclusion of more diverse real-world data, including more smaller-sized objects and overlays of real-world background images with synthetic objects, could help improve the performance of the mixed-data model. We expect that training with real-world background images, which have more textured skies due to weather conditions and camera optics compared to synthetic plain backgrounds, will help to decrease the number of false positives triggered by clouds, for example, and to increase the overall detection performance.

2.3.3. Benchmark for SORT on the synthetic video dataset

SORT [35] is a multi-object tracking algorithm which connects individual frame-by-frame object detections made by YOLOv5 into object trajectories. It relies on a Kalman filter and the Hungarian algorithm. The Kalman filter is used in SORT to predict the future location of objects based on their past states (e.g., position, kinematics), under the assumption of constant apparent velocity in the image. The covariances in the Kalman filter represent the uncertainty associated with the predicted state. The Hungarian algorithm, on

the other hand, is used to solve the assignment problem, where the goal is to associate new detections with existing tracks. The algorithm takes in a cost matrix, where each element represents the “cost” or mismatch distance (in our case, using the IoU) of associating a detection with a track. The Hungarian algorithm then finds the optimal matching that minimizes the total assignment cost, ensuring that each detection is paired with the most appropriate existing track.

We analyze the performance of our SORT implementation on the synthetic video dataset, which has three different types of trajectories: straight, curved, and piecewise, as shown in Figure 10. We set a IoU threshold of 0.5 to define a match between bounding boxes predicted by the Kalman filter and the ground truth bounding boxes. Each reconstructed trajectory is assigned an identifier (ID). We report the following standard metrics in addition to the ones defined in Section 2.3.2, for tracking performance [39], [40], [41], [42]:

- *Identity Switches (IDS)*, also known as association error, the count of how many times the reconstructed ID associated with a ground truth trajectory changes, i.e. where objects were incorrectly re-identified as new objects
- *Multiple Object Tracking Accuracy (MOTA)*: a metric compiling tracking errors over time, $MOTA = 1 - \frac{\sum_t FN_t + FP_t + IDS_t}{\sum_t TP_t + FN_t}$
- *Multiple Object Tracking Precision (MOTP)*: a measure of localization accuracy, $MOTP = \frac{\sum_t d_{t,i}}{\sum_t c_t}$ where c_t is the number of matches in frame t , and $d_{t,i}$ is the bounding box overlap of target i with its corresponding ground truth box
- Track quality metrics relative to ground truth trajectories that have been successfully tracked, i.e. have matched Kalman filter predictions, regardless of their predicted ID:
 - *Mostly Tracked (MT)*: for at least 80% of their life
 - *Mostly Lost (ML)*: for less than 20% of their life
 - *Partially Tracked (PT)*: for between 20% and 80% of their life
 - *Track Fragmentations (FM)*: a count of how many times a ground truth trajectory goes from tracked to untracked status
- Identification metrics establishing a one-to-one match between ground truth trajectories and reconstructed trajectories:
 - *ID Precision (IDP)*: the fraction of reconstructed trajectories which have a match
 - *ID Recall (IDR)*: the fraction of true trajectories which have a match
 - *ID F1 Score (IDF1)*: the harmonic mean of IDP and IDR

SORT parameters optimization

Our implementation of SORT uses three standard parameters:

- *max_age* = 100 frames, which determines how long a reconstructed track can exist without a match before being deleted from the list of current reconstructed track
- *min_hits* = 3, which requires a track to have a minimum number of consecutive matches before being confirmed
- *iou_threshold* = 0.3, which sets the minimum IoU needed for a detection to be associated with a track

and a custom parameter to mitigate the intermittent lack of YOLOv5 detections:

- *scale_factor* = 4, which is a multiplicative factor to scale all bounding boxes width and height proportionally before running SORT

We confirm that the values we used for *iou_threshold* and *scale_factor* are optimal or close to optimal, balancing tracking accuracy and efficiency in our specific application.

First, we tuned the *iou_threshold* by varying it from 0.01 to 0.90 and computing both tracking errors and track quality metrics. The right of Figure 12 shows the number of identity switches (IDS) divided by the total number of ground truth trajectories, across different IoU thresholds. As the threshold increases from 0.01 to 0.90, a rise in identity switches was observed, particularly at IoU thresholds of 0.60 and 0.90, where the normalized IDS peaks

at 0.06 and 0.24, respectively. Lower thresholds (≤ 0.30) exhibit far fewer identity switches, indicating fewer instances where objects were incorrectly re-identified as new objects.




Figure 12. Left: Track quality metrics MT, PT and ML as a function of IoU threshold; values are normalized by the total number of ground truth trajectories (209) and together add up to 1. Higher values of MT (mostly tracked) are better. Right: IDS normalized by the total number of ground truth trajectories (209) as a function of IoU threshold. Lower values of IDS (track identity switches) are better. .

The relationships between IoU thresholds and the tracking quality metrics MT, PT, and ML are shown in the left side of Figure 12. The number of ground truth trajectories remains constant at 209. The MT fraction is highest at low IoU thresholds (0.01 - 0.30) and worsens as the threshold increases. At thresholds ≥ 0.4 , MT begins to drop monotonically, indicating that at higher IoU thresholds, fewer ground truth objects are being correctly tracked throughout their entire trajectory. The PT fraction remains relatively stable across all thresholds, with values fluctuating between 0.06 and 0.10, suggesting that partial tracking performance is not influenced by the IoU threshold. In contrast, the ML fraction increases monotonically with higher IoU thresholds, from a minimum of 0.07 at an IoU of 0.01, and to a maximum of 0.55 at an IoU of 0.90. The relationship between IoU thresholds and IDS (reconstructed trajectory identity switching) is shown on the right side of Figure 12. The abrupt increase in IDS between IoU thresholds of 0.50 and 0.60 indicates a sudden jump in the proportion of trajectories where objects were incorrectly re-identified as new object.

The results of these two track quality analyses suggests that an IoU threshold around 0.3 provides the optimal balance between minimizing identity switches (low IDS) and maintaining accurate tracking (high MT), as it results in normalized IDS of 0.03, MT of 0.79, ML of 0.12. In contrast, higher IoU thresholds (≥ 0.8) lead to excessive re-identification errors, while lower IoU thresholds (≤ 0.2) risk over-tracking due to lenient matching criteria between the Kalman filter prediction and the ground truth bounding box. In all following studies, we used $iou_threshold = 0.3$.

In our implementation of SORT, instead of tuning the default parameters of the Kalman filter, we artificially increase the bounding box (BBox) size by multiplying both the height and width by a scale factor. This adjustment addresses the challenge of tracking objects with varying bounding box size, such as objects moving towards or away from the camera. Initially we picked a scale factor of 4 based on visual inspection, but the synthetic video dataset enables us to check that this choice is close to optimal. Figure 13 shows the MOTA, IDF1 and MOTP metrics as a function of the BBox scale factor which ranges from 0.5 to 30.0. MOTA in the middle of Figure 13 improves rapidly at first, peaking around a BBox scale factor of 5-10. The IDF1 graph on the left measures the tradeoff between identity precision and recall for tracking. We see a sharp rise in performance initially, with a peak around a scale factor of 5-10. Beyond a scale factor of 10, the IDF1 metric declines. The MOTP graph on the right shows that as the BBox scale factor increases, the MOTP decreases. The smallest MOTP values are observed at the scale factor of 10-20. This confirms that our original, heuristic choice for the BBox scale factor of 4 is not too far from optimal. Overall it




Figure 13. MOTA, IDF1 and MOTP tracking performance metrics for our SORT implementation for three types of synthetic trajectories, as a function of the BBox scale factor .

results in a 40% increase in the number of correctly matched tracks, compared to a BBox scale factor of 1. For all the following studies we use a BBox scale factor equal to 4.

Benchmark on the synthetic video dataset

With these parameters in our SORT implementation, we also analyzed performance on the synthetic trajectories as a function of trajectory type: straight, curvy, and piecewise. The advantage of synthetic trajectories is that we control and can isolate the trajectory characteristics influencing the tracking performance.

Table 2. SORT tracking performance metrics for different trajectory types from the synthetic video dataset.

Metric	Curved	Piecewise	Straight	All
FP count	1757	1053	800	3610
FN count	9871	5208	3258	18337
Recall	87%	92%	92%	90%
Precision	97%	98%	98%	98%
IDs count	122	149	15	286
MOTA	0.84	0.90	0.90	0.87
MOTP	0.14	0.13	0.14	0.14
MT / PT / ML	984 / 174 / 206	743 / 104 / 74	479 / 65 / 60	2206 / 343 / 340
FM count	380	353	202	935
IDP	95%	94%	97%	95%
IDR	85%	88%	91%	87%
IDF1	89%	91%	94%	91%

The results in Table 2 show that across all three types of trajectories, we find a MOTA of 0.87. It also demonstrates that SORT exhibits strong identification precision (IDP = 95%) and high object recall (Rcll = 90%) across synthetic data, while maintaining a high tracking accuracy with MOTA = 0.87. However, a moderate level of fragmentation (FM = 935) suggests that identity switches and interruptions occurred during the tracking of objects, likely due to the challenges posed by rapid or unpredictable object movements in the synthetic environment specifically for curved trajectories. The near-perfect precision (Prcn = 98%) indicates that the model had very few false positives. We investigate further the tracking performance for three types of trajectories, and as seen in Table 2, the model's performance varies depending on the nature of the object trajectories. For the 604 straight trajectories, the model achieved its highest overall performance, reflected by an IDF1 score of 94% and a MOTA of 0.90, due to the simplicity of tracking linear motion. Precision and recall metrics for straight trajectories are also high, with Prcn = 98% and Rcll = 92%, respectively. These results suggest that the model efficiently tracks objects following predictable, straightforward paths, minimizing identity switches (IDs = 15) and fragmentation (FM = 202). For piecewise trajectories, which involve sudden changes in object direction, the algorithm performs better than on curvy trajectories, with an IDF1 of 91% and MOTA of 0.90. The slight improvement compared to curvy trajectories is likely due to the model's ability to adjust more easily to segmented motion changes rather than continuous curving




Figure 14. Left: YOLO detection count per camera, first divided by the camera’s visible sky area and total recordings count, and then normalized across cameras so all values sum up to 1. This histogram includes all detections of any object from a subset of the five months of commissioning data. Camera 8 was offline during this interval. Middle and Right: 2D histograms showing examples of spatial distribution of YOLO detections with confidence score > 0.5 , for the month of May 2024, for camera 5 (SE) and 7 (SW), respectively. A regular commercial aircraft route is clearly visible in camera 5 (center).

movements. However, the number of identity switches (IDs = 149) and fragmentation events (FM = 353) is noticeably higher for piecewise trajectories, reflecting the difficulty the model faces in consistently associating object identities across sudden directional changes.

3. Commissioning results

In this section we focus on the Dalek as a whole system: its eight IR cameras and the pipeline from data to reconstruction of aerial events. We present basic commissioning checks on the recorded dataset and corresponding reconstructed quantities such as individual frame-by-frame detections or object trajectories. After these, we characterize more in-depth the performance of the Dalek and the detection pipeline as a function of atmospheric conditions and object characteristics using two distinct datasets: ADS-B derived dataset, which focuses on airplanes in real-world data, and a synthetic dataset.

3.1. Basic checks on recorded dataset

3.1.1. Environment effects on spatial distribution of detection counts per camera

We start with the detection counts per camera, and attempt to correlate any long-term spatial inhomogeneities with known conditions at the observatory site. We check each camera for its spatial distribution of YOLO detections to rule out any potential bias. ADS-B historical records confirm that cameras 5 and 7 capture airplane high traffic routes. This is shown in the center of Figure 14, where a commercial airline approach route to a major airport is visible when plotting purely detection counts. We also find that Cameras 1 and 4 often capture the moon at night (not shown) which triggers a detection by YOLOv5. Additionally, detections caused by birds’ preferential perches or nests have small spatial signatures at the edges of the treeline. New leaves and growth of the trees beyond the treeline image mask that we utilize may trigger an unusually high count of detections in certain months. Some spurious points may be due to dust on the lens of the Boson camera or imperfections in the coating of the Germanium windows. Blowing leaves and the edges of clouds are a source of many detections, but we expect them to be randomly distributed over time scales of months. And finally, the cameras regularly display a green square in the top right corner of the image during the automatic flat-field calibration, which can be detected by YOLO and show up as an unusually high bin in this histogram, clearly visible in Camera 5, Figure 14 (center image).

3.1.2. Cross-camera check

We check cross-camera variations of the detection counts by normalizing the detection count by area of the sky above the treeline, and by total recorded time: some cameras




Figure 15. Left: average hourly detection count for different cameras throughout the day. Middle: 2D histogram of the detections’ bounding box width and height showing pronounced clustering that may be useful for object classification. Right: 2D histogram of the detections’ bounding box area and aspect ratio, again with strong clustering. The histograms include all detections of any object from a subset of the five months of commissioning data.

pointing north are constantly in recording mode with their sunshades open, while others alternate between recording and pause mode, sunshade open and closed, depending on the sun position in the sky. Figure 14 shows on the left that the average normalized detection count is not uniform across cameras. We do expect that air traffic routes or the Moon, mentioned previously, can bias this metric for certain cameras.

3.1.3. Bounding boxes properties

The middle of Figure 15 is a histogram of YOLOv5’s detection bounding box width versus height and highlights a potentially useful feature where distinctly different populations of detections appear as elongated clusters. The histogram includes represent all detections, not just airplanes, and currently we can only speculate about which object falls into which cluster. For example, the population cluster close to a bounding box aspect ratio of 1 would correlate with the regular Moon detections. Clustering also appears in the right-hand side of Figure 15 which compares bounding box aspect ratio with its area. The variety in bounding box size and aspect ratio suggests that they may be useful features in the future to help classify objects and search for outlier populations.

After these basic checks on the Dalek detections, we studied more in-depth the performance of the Dalek and the detection pipeline, including as a function of atmospheric conditions and object characteristics, on two distinct datasets: ADS-B derived dataset, which focuses on airplanes in real-world data, and a synthetic dataset.

3.2. Performance evaluations using ADS-B equipped aircraft

In this section we evaluate the physical performance envelope of the full Dalek system using data related to the many ADS-B equipped airplanes (Section 2.2.3) that fly over our development site. We use the latitude, longitude, and altitude from an ADS-B derived real-world dataset to be the aircraft’s true positions, project these onto each 2D image frame, and compare these true points to the objects predicted by YOLOv5’s detection algorithm.

3.2.1. Evaluation methodology

For this analysis we define sets of ADS-B entries using the following sequential selection criteria:

1. Within a square of side 10 km centered on the observatory,
2. within the field-of-view of at least one camera,
3. above the treeline of the camera from which they should be visible,
4. at a time when there is a recording by the relevant camera, and finally,
5. are detected by YOLOv5

For each selected ADS-B record, we use the camera’s extrinsic and intrinsic calibration to translate the 3D coordinates of the aircraft to 2D coordinates in the frame of the camera images. We match ADS-B records with existing bounding box detections using a temporal

threshold of 5 seconds and a spatial threshold of 30 pixels. These thresholds are generous to account for the imprecision of the recording timestamps in the time period we selected and the potential error in the reported ADS-B positions.

To define performance metrics, we first identify the number of aircraft within range of our site (criteria 1) and define these as *in-range*. From this set, we identify those that were in the right location to have been viewed by the Dalek, that is, within the effective field of view of at least one camera (criteria 2 and 3). The set of aircraft that meet the first three criteria outlined above is defined as *viewable*, and it is independent of airplane size, aspect, distance (under 10 km), ambient IR conditions, cloud cover, and camera up-time. We then identify the subset of these viewable aircraft who were also in the camera's recording window (criteria 4), that is, they were in the right place and at the right time. The set of aircraft that meet the first four criteria is defined as *recorded*.

The *acceptance* is related to instrument up-time and is defined here as the number of *recorded* aircraft divided by the number of *viewable* aircraft. In other words it is the fraction of viewable aircrafts that the Dalek managed to record; these aircraft were "accepted" into the analysis pipeline. The set of image frames with recorded ADS-B aircraft is then compared with YOLOv5 detections for the same frames (criteria 5). Our final set consists of aircraft matching the time and place of objects detected by YOLOv5, meeting all five criteria. We define this set as *detected*.

The *efficiency* is related to YOLOv5 performance, and is defined here as the number of *detected* aircraft divided by the number of *recorded* aircraft. In other words, it is the fraction of aircrafts meeting the first four selection criteria above which are also detected by the YOLOv5 object detection stage. If our pipeline is efficient, it converts most of the objects that enter the pipeline into detections, with a minimum of loss. This metric can depend on aircraft size, aspect, distance, and atmospheric conditions. If our Dalek system pipeline and site location were perfect, all aircraft within our site range limit would be in our *detected* set. Our pipeline is not perfect, and these sequential metrics help us find areas to focus improvements. We now examine these metrics.

3.2.2. Performance results

Overall, we counted 27,467 airplanes who flew within a radius of 10 km centered on the observatory, of which 8,550 met the four selection criteria outlined above, and, of these, 3,678 were matched in at least one frame to a YOLOv5 detection bounding box, meeting criteria 1-5. We use this number to estimate the fraction of all objects reconstructed during the commissioning period that is made up of aircrafts. Taking the total count of reconstructed object tracks from Section 3.4.1, which includes aircraft, birds, leaves, clouds, etc., we estimate that $\sim 0.7\%$ of all reconstructed object tracks were associated with ADS-B equipped aircraft.

We use the five months of data from the commissioning period to examine our ability to reconstruct aerial objects such as aircrafts from their physical parameters using only our Dalek system. We look at the sequence of processed information, from all the aircraft that passed overhead, to the detections that come out of YOLOv5. We can quantify the falloff from the average daily count of ADS-B airplanes within range of our site, to those also within the effective field of view of at least one camera, to those also within that camera's recording time window and, finally, to those also detected by YOLOv5. For example, the upper left panel of Figure 16 shows this progression in average daily counts. For a given aircraft, we bin its ADS-B records according to their distance to the site, before counting unique aircrafts in each distance bin. For example, we expect on average, within ~ 10 km of the site, ~ 210 airplanes to be viewable in a day, of which ~ 76 pass by within a recording window and ~ 31 are detected. At a distance of 10 km, the mean efficiency is 43%, the rest of the recorded airplanes at this distance were not detected by YOLOv5. The detection efficiency decreases with increasing distance of the airplane to the observatory. We also see from examining the ADS-B data that there are flight lanes to regional airports




Figure 16. Upper left: For the commissioning period, average daily count vs distance (km) from observatory of: ADS-B equipped airplanes in range of site (criteria 1: *in-range*); of those, also within the effective field of view of at least one camera (criteria 1-3: *viewable*); of those, also within that camera’s recording time window (criteria 1-4: *recorded*); and, of those, also detected by YOLOv5 (criteria 1-5: *detected*). Middle left: the number *detected* divided by the number *recorded* (*efficiency*), vs distance (km). Lower left: the number *recorded* divided by the number *viewable* (*acceptance*), vs distance (km). Right: *acceptance* and *efficiency* as functions of apparent airplane size. The histogram shows the number of ADS-B records that contributed to each point on the graph. Error bars are computed by propagating statistical errors from all ADS-B counts, assuming Poisson distributions.

within 2-5km horizontal distance to the observatory, which, accounting for typical aircraft altitudes, correlates with the peak of aircrafts in-range around ~ 5 -7km.

On the right of Figure 16, we examine the effect of pixel size on both acceptance and efficiency. The aircraft’s apparent size in pixels can be inferred from the track angle (heading) reported in most ADS-B records, the 3D location of the aircraft with respect to the observatory, the true aircraft size and the camera sensor field of view. We take the maximum between the projected width and projected length. If the track angle is not available, we assume that the aircraft track is orthogonal to the line of sight.

We see that acceptance varies slightly with apparent size in pixels. It is counter-intuitive for there to be a relationship between the *recorded* set and pixel size, unless airplane apparent pixel size is related to field of view or recording time intervals. Recording intervals and aircraft flight schedules both depend on the time of day, and aircraft sizes are related to their schedules. Large aircraft, on regular commercial flights on approach in the evening, small to mid-sized private jets arriving at the end of the day, and small propeller planes operated by a nearby flight school at mid-day, all have their own distinctive flight paths and schedules. For example, a flight path used by large commercial jets on approach to a major airport, can be seen in the center of Figure 14, with many of these flights arriving in the early evening, when they are captured and detected in Camera 5 (faces south-east with sunshade open). The correlation between aircraft size and time of day is shown in the center of Figure 17. We see from the efficiency curve in the right of Figure 16 that YOLOv5 detections, on the other hand, are least efficient at 1 pixel, increase with pixel size until a maximum at five pixels, and the drop again until 10 pixels. This matches the known behavior of YOLOv5 related to pixel size, as seen in trials using synthetic data.

In Figure 17 we examine actual size of aircraft and its effect on acceptance and efficiency in the Dalek dataset, using length as a proxy for size. Again counter-intuitively, the acceptance does depend on the airplane real size. We speculate this is caused by a bias in the time of day that certain size aircraft fly. The 2D histogram in the the middle of Figure 17




Figure 17. Left: efficiency as a function of actual airplane size. Error bars are computed by propagating statistical errors from all ADS-B counts, assuming Poisson distributions. Middle: actual airplane size distribution throughout the day. Right: actual airplane size distribution compared to projected size in the image.

shows that small airplanes, many of which are associated with a nearby flight school, are more likely than large airplanes to fly in the middle of the day, when we have more camera shades closed to avoid exposure of the sensors to the sun. The efficiency, on the other hand, decreases when the airplane real size increases, which is also counter-intuitive at first. The right plot in the same figure shows that, on average, small and large airplane length correlate with larger apparent size, which is likely a consequence of the specific air traffic routes that are visible from our site.

Using weather data that is correlated with ADS-B records (see Section 2.3.1.5), we examine the effect of precipitation, visibility, relative humidity, and temperature on the acceptance (recording) and efficiency (detection) of airplanes. As seen in Figure 18, precipitation, visibility, and relative humidity have only a slight impact on acceptance, as might be expected. The effect of temperature on acceptance, shown in Figure 19, is more pronounced than the other weather effects, with acceptance decreasing monotonically as temperature increases. We correlate this unexpected relationship to the fact that daily temperature cycles are related to the sun's daily cycle, which ties temperature to the sunshade and recording cycle, which in turn, affects acceptance. The same figures show that increased precipitation, decreased visibility, and increased humidity have a marked adverse effect overall on efficiency. This is consistent with how moisture affects IR atmospheric transmission and therefore object image quality: we know that the atmosphere's relative humidity impacts transmission in the infrared wavelengths. Temperature mildly affects efficiency, as shown in Figure 19.

We then look at efficiency as a function of camera and pixel location within the image frame, in Figure 20. Camera 8 has a wider FOV and more fish-eye lens distortion compared to the other seven cameras, and we see that the detection efficiency in camera 8 is lower than most cameras, with the exception of camera 2. Camera 2's extremely low efficiency is primarily due to its low uptime and thus low acceptance until a damaged USB cable




Figure 18. Left: efficiency as a function of precipitation. Middle: efficiency as a function of visibility. Right: efficiency as a function of relative humidity. Error bars are computed by propagating statistical errors from all ADS-B counts, assuming Poisson distributions.

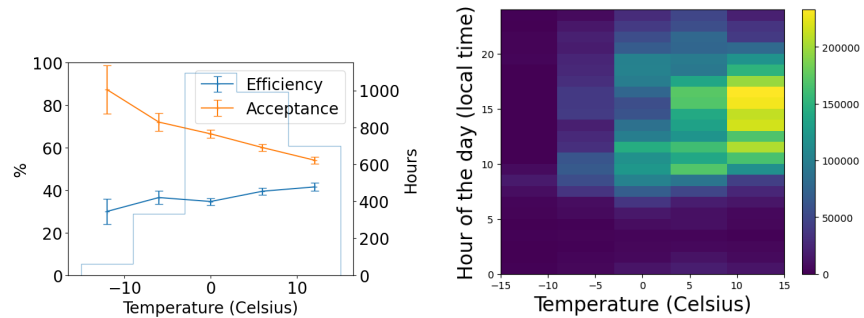


Figure 19. Left: efficiency as a function of temperature. Right: uneven temperature distribution throughout the day in our dataset. Error bars are computed by propagating statistical errors from all ADS-B counts, assuming Poisson distributions.




Figure 20. Left: efficiency as a function of camera. Middle: efficiency as a function of horizontal position. Zero corresponds to the left of the image. Right: efficiency as a function of vertical position. Zero corresponds to the top of the image. Error bars are computed by propagating statistical errors from all ADS-B counts, assuming Poisson distributions.

was diagnosed. The dip in efficiency around a horizontal position of 400 px is related to a radio antenna tower in camera 1 whose silhouette is included in the treeline mask: since detection bounding boxes which overlap with the antenna are removed from this analysis, an airplane passing by the tower may not have matching detections in the tower area even if it is successfully detected elsewhere in the image. The increase in efficiency close to the center of the image on a vertical axis coincides with the average vertical position of the treeline, which may artificially inflate the efficiency: branch movement and tree growth over time trigger more false positives for the YOLOv5 detection algorithm, which may accidentally be matched to a passing airplane which is not detected elsewhere in the image.

Overall the Dalek frame-by-frame object detection stage has an acceptance of 41% for ADS-B equipped aircrafts within a square of side 10 km centered on the observatory, and a mean detection efficiency of 36%. Acceptance is heavily dependent on range, field of view, and camera uptime. Efficiency depends very sensitively on weather conditions such as precipitation, atmospheric visibility and relative humidity, as well as the condition of the lens, dust and raindrops. It is also not perfectly uniform over camera images due to boundary effects of the detection algorithm and interference from treeline growth. Figure 21 summarizes the fraction of aircraft *in-range* which are *detected*, per camera and spatial location within the camera image frame. This visualization convolves the effects of recording schedule, downtime, weather conditions, detection efficiency and underlying aircraft statistics, such as flight lanes from the BOS airport.

3.3. Evaluation using a synthetic dataset

While ADS-B records allow us to evaluate the performance of the full pipeline on real-world data, airplanes, which have predictable trajectories and a limited range of kinematics and characteristics, are not the only objects that we detect using the Dalek. To estimate how reconstruction performance might extend to known objects with different kinematics and to quantify this based on the object trajectory properties, we generate a synthetic video

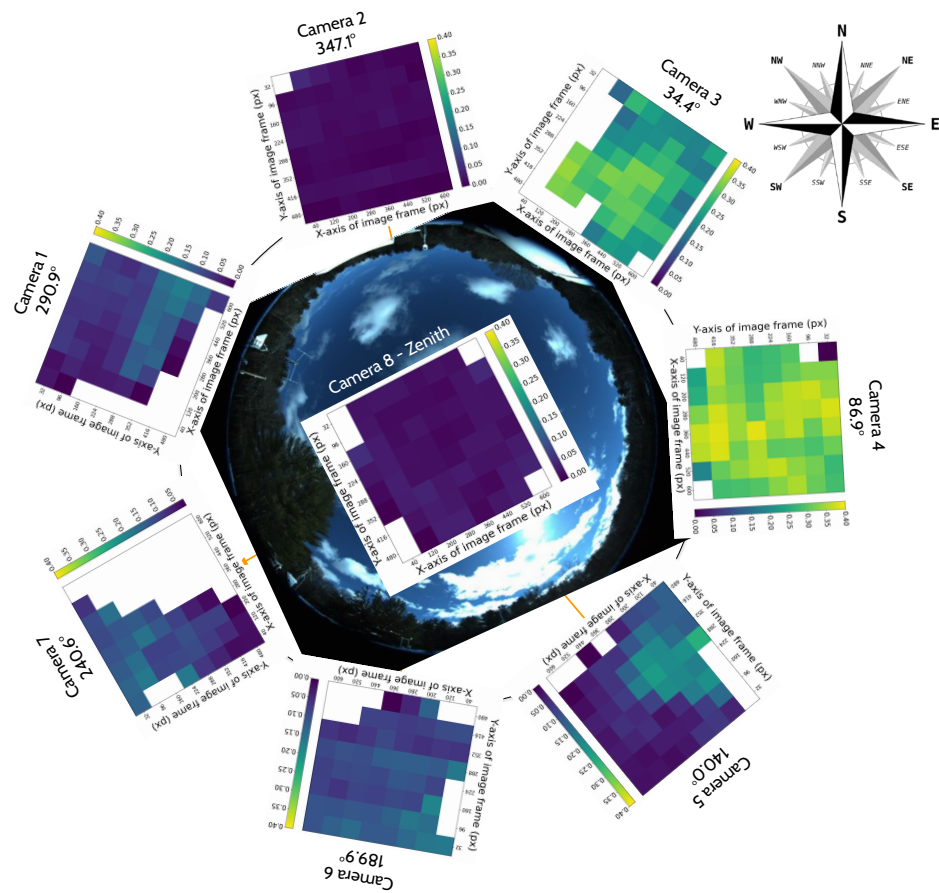


Figure 21. Fraction of aircraft *in-range* which are *detected*, for each camera and spatial location in the camera image frame. Each bin represents an area of size 80x64px in the original camera image. The origin of both axis is set to the upper left corner, following the computer vision convention. White bins are either entirely included in the treeline mask, where we do not look for detections, or happen to lack any aircraft *in-range* during the commissioning period.

dataset using a variety of objects exhibiting straight, curved, and piecewise trajectories, as described in Section 2.3.1.3.

We first run YOLOv5 for object detection on the individual frames of the videos, and keep only bounding boxes larger than 3x3 pixels from both true (labels) and predicted (detections) bounding boxes. We consider it a match if a true and predicted bounding box have an intersection over union (IoU) greater than 10%. 14% of the true trajectories do not contain any points with a bounding box matched to a YOLOv5 detection.

In the next step, we run SORT on YOLOv5's output to form reconstructed trajectories from the frame-by-frame detections. We find that 27% of object detections, although matched to a true bounding box of an object that has a true trajectory, are not then associated with any reconstructed trajectory. Conversely, we are also interested in the fraction of true trajectories which contain at least one matched bounding box, but where these individual detections are not reconstructed into a predicted trajectory. We find that 8% of the true trajectories in this dataset are not reconstructed even though SORT was given matched detections from YOLOv5. Table 3 breaks down these metrics for each of the three types of trajectory.

We compute the efficiency or fraction of true trajectories which are matched to at least one reconstructed trajectory, as well as the purity or fraction of reconstructed trajectories which match with a true trajectory. A match between a true and a reconstructed trajectory

Trajectory type	Curved	Straight	Piecewise
Unique true trajectories	~430	~740	~440
Detection precision (<i>fraction of detections which are matched to a true bounding box</i>)	50%	41%	51%
Detection recall (<i>fraction of true bounding boxes which are matched to a detection</i>)	92%	54%	91%
Fraction of matched object detections without an associated reconstructed trajectory	29%	29%	24%
Fraction of true trajectories which have at least one matched bounding box	99.8%	81%	99%
Fraction of true trajectories which have at least one matched bounding box but no associated reconstructed trajectory	7%	15%	5%
Efficiency (<i>fraction of true trajectories matched to at least one reconstructed trajectory</i>)	94%	74%	96%
Purity (<i>fraction of reconstructed trajectories matched to a true trajectory</i>)	90%	83%	90%

Table 3. Breakdown of object detection and trajectory reconstruction metrics for different types of trajectories in the synthetic dataset.




Figure 22. Left: 2D histogram of reconstructed and true apparent speed. Middle: purity as a function of reconstructed apparent speed. Right: efficiency as a function of true apparent speed. Error bars are statistical uncertainties reflecting the bin population count.

is declared when at least 3 individual bounding boxes from the reconstructed trajectory are matched to true object bounding boxes, i.e. with an IoU greater than 10%. This means that a true trajectory can be matched to multiple reconstructed trajectories, if it was fragmented during the reconstruction. Efficiency and purity metrics are shown in Table 3 per trajectory type. Overall, we find an efficiency of 81% and a purity of 87%.

Since we lacked distance estimation during this commissioning period, distance, size, velocity, and acceleration of an object are degenerate quantities. Instead, we focus on the apparent area of the bounding box and the apparent projected 2D speed of its center. Figure 22 examines the trajectory efficiency and purity as a function of apparent speed, which is uniform along each trajectory in this dataset. As seen in the left of Figure 22, a very low (<1 deg / sec) reconstructed apparent speed seems to characterize pathological trajectories, since they can correspond to a wide range of true apparent speeds. In the middle and right of Figure 22, purity and efficiency have no obvious correlation to reconstructed and true apparent speed. The mean efficiency trends down for higher true apparent speed, for which we do have fewer true points to detect and reconstruct a trajectory, but smaller uncertainties are required to confirm this speculation.

Using the same metrics, Figure 23 examines the trajectory efficiency and purity as a function of apparent area in degrees squared. Purity shows variations beyond the statistical error bars between different bins of reconstructed apparent area, which may suggest that false positives are unevenly distributed among different bounding box sizes (apparent area) and certain aspect ratios may be more prevalent. Efficiency shows a slight improvement when the true apparent area increases.

Figure 24 focuses on the curved and piece-wise trajectories. The left most plot shows that for curved trajectories, a large curvature correlates with lower purity, i.e. there are fewer true positives. The efficiency (not shown) has no discernable correlation with the true curvature for curved trajectories. For the piecewise trajectories, one might intuitively expect the efficiency to drop when the inflection points counts increases, but the right side




Figure 23. Left: 2D histogram of reconstructed and true apparent area. Middle: purity as a function of reconstructed apparent area. Right: efficiency as a function of true apparent area. Error bars are statistical uncertainties reflecting the bin population count.




Figure 24. Left: Purity as a function of reconstructed curvature for curved trajectories. Right: Efficiency as a function of inflexion point counts for piecewise trajectories. Error bars are statistical uncertainties reflecting the bin population count.

of Figure 24 suggests that there is no significant variation: only the uncertainty of the metric increases. We attribute this to the non-unique trajectory matching scheme, which allows multiple reconstructed trajectories to be matched to the same true trajectory. In other words, the data presented in this figure do not account for trajectory fragmentation. Similarly, the purity (not shown) shows no significant variation as a function of the reconstructed apparent inflexion points.

Figure 25 examines how much fragmentation the reconstructed trajectories display. On the left is shown the number of synthetic videos having the true and reconstructed trajectory count shown on the x-axis. There are approximately three times as many reconstructed trajectories as true ones. This indicates that, on average, true trajectories are broken into ~ 3 fragments after passing through the pipeline YOLOv5 + SORT. On the right of the same figure, we compare the number of points in true versus reconstructed trajectories. A given true trajectory can have at most 100 points (since the synthetic videos in this dataset have at most 100 frames), which is why the distribution of true trajectory point count peaks at 100. The reconstructed trajectories are split in two distributions: some are successfully complete at 100 points, and some are presumably broken into fragments, each with a peak point count around 10.

In summary, among the true trajectories generated for this synthetic dataset, about 15% lack either frame-level object detection or trajectory reconstruction and are thus entirely missed. About 13% of the reconstructed trajectories are spurious and do not correspond to any true trajectory. The reconstruction of apparent speed is severely hampered by the fragmentation of trajectories, which causes a true trajectory to be broken on average into 3 smaller fragments of trajectory.

3.4. Unimodal aerial census

[4] described the key steps to search for scientific anomalies in the sky. The first step is a census of all aerial phenomena, whether man-made or natural. Although the Dalek

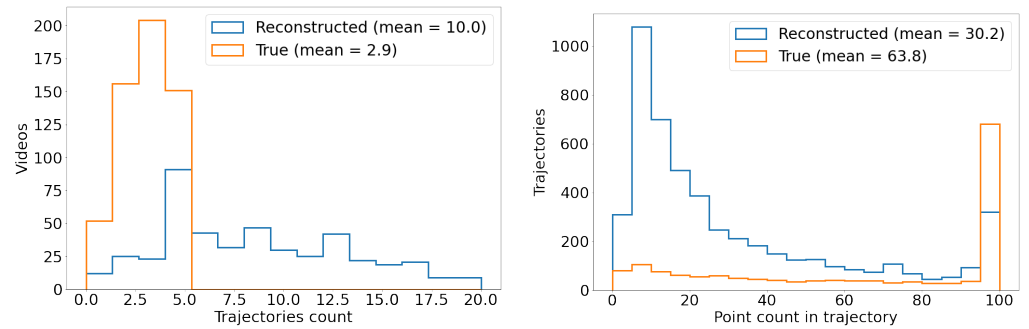


Figure 25. Left: Distribution of individual trajectory counts per video. The reconstructed trajectories count is higher on average due to trajectory fragmentation. Right: Distribution of trajectory point counts. True trajectories can have at most 100 points due to the dataset generation parameters, and reconstructed trajectories have fewer points on average along the trajectory, compared to true trajectories, due to missed or dropped detections.

data collected during this commissioning period was not yet science-quality, only involved a single instrument and did not have access yet to distance estimation for kinematics measurements, we look at the objects reconstructed by our pipeline in order to pave the way for the coming full-scale and multi-modal aerial census. We conduct a deliberately simplistic or *toy* search for outliers, in order to illustrate how we will apply a likelihood statistical analysis to any upcoming full-scale aerial census.

3.4.1. Toy outlier search

Dataset: reconstructed trajectories

For this toy outlier search, we look only at the reconstructed trajectories from the five months of commissioning the Dalek. The search is only conducted on the hemispherical seven cameras of the Dalek, excluding recordings from the zenith camera (camera 8). First, we exclude trajectories with fewer than four recorded data points, which do not provide sufficient information for apparent trajectory analysis. Secondly, we filter out the data points from reconstructed trajectories that intersect with the masked area of each camera, i.e. the area at or below the treeline, eliminating potential excessive detections caused by distant objects or environmental noise such as a growing treeline, leaves or birds perching in the treeline. After the masking and initial cleaning, we fit each reconstructed trajectory with cubic splines to reduce noise and improve the accuracy of calculated trajectory parameters. Between January and May 2024, from the seven hemispherical cameras, we collected and reconstructed 502,015 trajectories.

In Figure 26, we show examples of real-world reconstructed trajectory data points. Each color represents a reconstructed object’s trajectory, and the numbers next to each trajectory are unique identifiers generated by the tracking algorithm (SORT). A true trajectory may be broken down into several reconstructed trajectories, each with their own unique identifier and color. Each data point is separated in time by at least one frame interval, i.e. 0.1 second. The distance between data points reflects the apparent reconstructed kinematics of the object and the performance of the frame-by-frame YOLO detection stage. The red and orange lines (labeled 118 and 15) in Figure 26 belong to an object moving horizontally across the image. The trajectories contain segments with uninterrupted detections, indicated by the high spatial density of data points, as well as large empty gaps: these objects likely display a steady, linear motion disrupted by an occasionally low detection rate. The green and blue trajectories (labeled 36, 124, 160, 50) are curvy and include gradual changes in direction. Looping and curvy trajectories, such as the purple (labeled 18) and the brown (labeled 153) ones, suggest more complex or erratic motions, which could be generated by birds, for example. With the SORT parameters that we tuned in Section 2.3.3, we are able to reconstruct some trajectories where an occlusion happened, such as those labeled 36 and 46, which intersect each other. Figure 26 also includes several long trajectories (labeled 46,

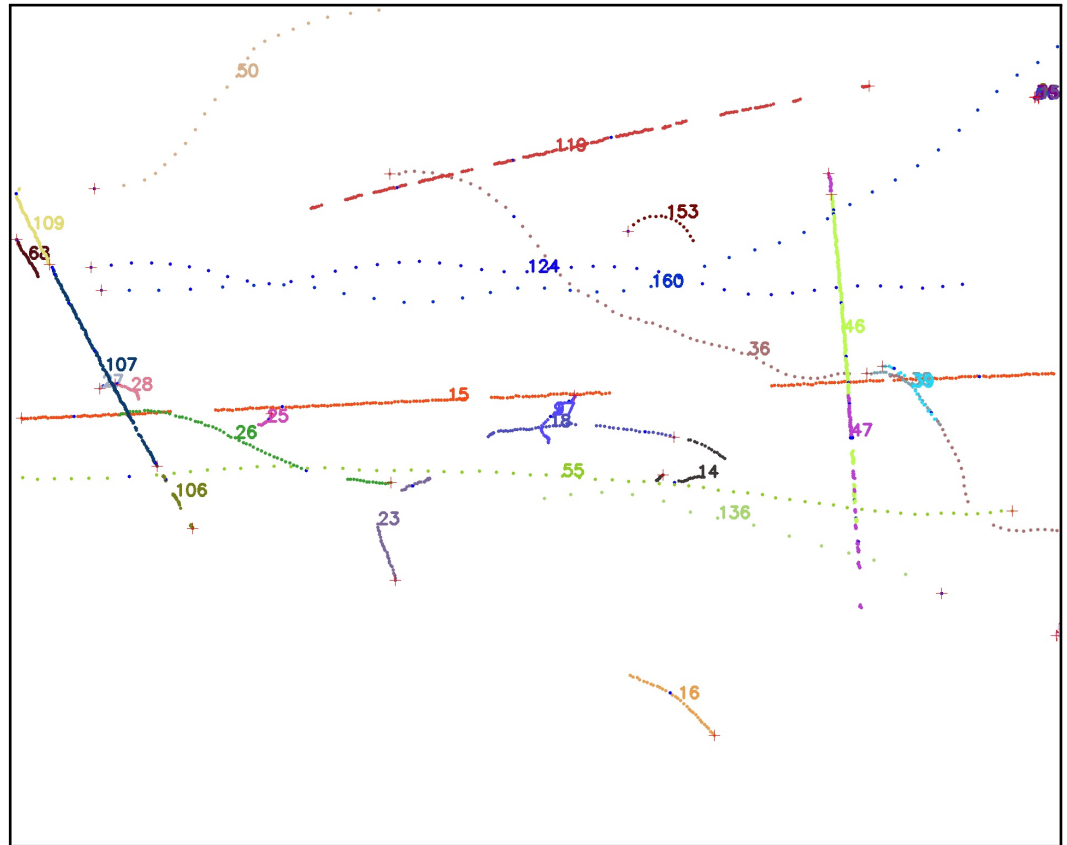


Figure 26. Overlay of multiple reconstructed trajectories, randomly sampled from different videos in the commissioning dataset. Each trajectory is plotted as a set of consecutive 2D points, representing the unique detections underlying the trajectory, and assigned a unique color. The numbers next to each trajectory are unique identifiers, assigned by the SORT algorithm.

47, 107 or 110) generated by an airplane at different times. Notably, trajectories 46 and 47 were produced by the same object. Due to the loss of detections for a significant number of frames, the tracking ID was switched, causing the trajectory to appear segmented. This illustrates the challenges posed by intermittent detection loss during tracking, resulting in the separation of a continuous path into multiple labeled segments.

Toy outlier criteria based on sinuosity

To flag outliers, we compute the *sinuosity* of each 2D reconstructed trajectory, which is defined as the total trajectory length divided by the length of the straight line from trajectory start to end points. Figure 27 shows the distribution of sinuosity values for randomly selected trajectories. We estimate the probability density of sinuosity values in order to identify high-sinuosity outliers, using both non-parametric and parametric approaches. Given the distribution's apparent non-Gaussian probability density, we compare Kernel Density Estimation (KDE) models with varying bandwidths, with a Gaussian distribution fit as a baseline. The distribution of sinuosity values is log-transformed to reduce skewness. We show in Figure 27 the fit of the different models to the underlying probability density. Bandwidth selection is critical in KDE as it controls the smoothing level; finer bandwidths capture greater details but risk overfitting, while larger bandwidths provide greater smoothness at the expense of capturing distribution nuances. We test KDE with different bandwidths, including Scott's and Silverman's methods as well as a custom fixed bandwidth of 0.1, to balance the smoothness and fit quality.

In Table 4 below, we compute the log-likelihood of each model's fit with the real distribution of sinuosity values. Higher (less negative) values indicate a better fit. KDE

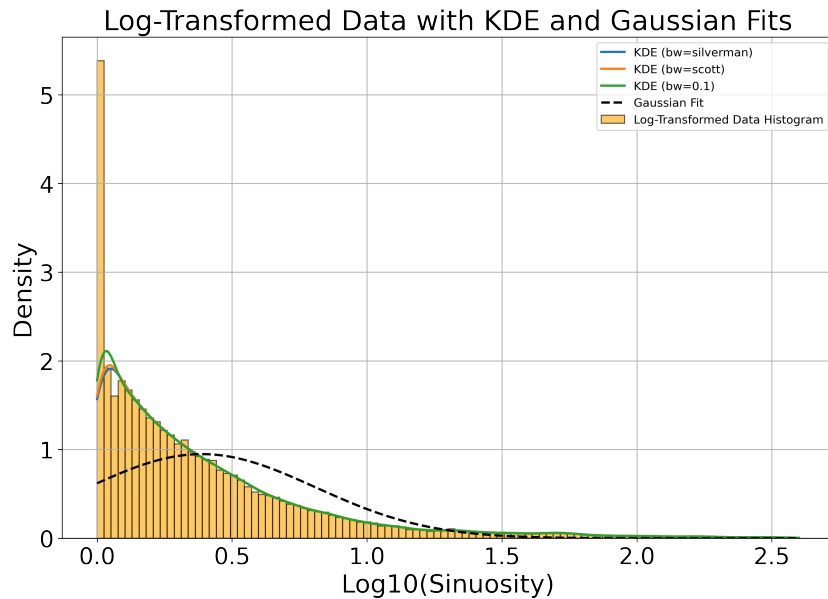


Figure 27. Normalized histogram of the log-transformed distribution of sinuosity values, computed for all reconstructed trajectories from the five months of commissioning. The lines show the fit with four models for probability density estimation: Kernel Density Estimation (KDE) with different bandwidths in colored solid lines and a Gaussian fit in dashed line.

with a 0.1 bandwidth yields the highest log-likelihood, outperforming both the other KDE with a different bandwidth and the Gaussian fit. The KDE model does best at capturing the underlying density of this dataset. The Gaussian fit's low log-likelihood suggests it was unable to accommodate the distribution's skewness.

Table 4. Log-Likelihood Comparison of KDE and Gaussian Models for the distribution of trajectory sinuosity values. A larger log-likelihood indicates a better fit. KDE with a fixed bandwidth of 0.1 performs best.

Model	Log-Likelihood on Test Data
KDE (bw = silverman)	-1097.19
KDE (bw = scott)	-1020.60
KDE (bw = 0.1)	-751.99
Gaussian	-6860.76

For this toy outlier search, we arbitrarily decide to use the 75th percentile as a selection threshold, which corresponds to sinuosity value of 3.0. The sinuosity distribution can be partitioned into three categories based on sinuosity ranges. Trajectories with a sinuosity between 1.0 and 1.2 tend to be straight and predominantly associated with planes and birds, with minimal contamination from clouds. Trajectories with a sinuosity between 1.2 and 3.0 are curvilinear. Trajectories exceeding a sinuosity of 3.0 are primarily attributed to clouds, with some contribution from birds and a negligible presence of planes. Clouds do not move very fast and they can appear almost stationary for a period of time: clouds' reconstructed trajectory tends to have a very short Euclidean distance (i.e. length of a straight line from start to end of the trajectory) which drives up the sinuosity value computed for these trajectories. Among the original dataset of 502,015 trajectories, we find that $\sim 16\%$ of them or 81,873 trajectories have sinuosity values greater than 3.0. This threshold makes it feasible for us to manually examine the distribution's upper tail, i.e. reconstructed trajectories with high sinuosity. For the purpose of this manual examination, to sift through the selected high-sinuosity trajectories, we generate two types of visualizations which are daily summaries of reconstructed trajectories. One is similar to Figure 26, explained previously. The other uses a background camera image computed from an averaged of 100

empty frames. We overlay on this background, for each trajectory, three snapshots insets in a green circle, which show the object's detection bounding box at the start, middle and end of the trajectory. This visualization reveals regular patterns of trajectories accumulated daily. Figure 28 illustrates the partition of sinuosity values and randomly sampled examples from each of the three ranges of sinuosity. The left and right columns display the two types of daily summaries we just described. The top, middle, and bottom rows of Figure 28 show examples of reconstructed trajectories with sinuosity in the ranges of $[1.0; 1.2]$ (predominantly identified as planes and birds), $[1.2; 3.0]$ (mostly birds and clouds) and $[3.0; \infty)$ (mostly clouds) respectively.

Manual examination of high-sinuosity trajectories

The approach described above helps to identify objects following unusually curved paths or exhibiting erratic movement patterns, distinct from the smoother motion of most tracked objects. It is followed by a manual review of the trajectories that passed the high sinuosity criteria. The majority of trajectories are associated with mundane objects. Figure 29 shows examples of reconstructed trajectories where the object can be identified by looking at the detection bounding boxes and the original recordings. Some can be identified as airplanes flying across the sky (e.g., trajectories 88, 21, 70, 19, 57 in Figure 29 (a)), others are identified as flocks of birds flying together in formations (e.g., trajectory 20 in Figure 29 (b), or trajectory 107 in Figure 29 (a)), rotorcraft (e.g., trajectory 162 in Figure 29 (c)), as well as clouds (14, 25 in Figure 29 (b)).

While some objects can be visually identified, especially when they are large in apparent size, many objects remain visually indistinct due to their smaller size. This is particularly evident from the bounding box sizes shown in Figure 15 (middle), which demonstrate that most detected objects are small. As a result, the majority of small objects remain unidentified through visual inspection alone. Despite having the reconstructed trajectory, the small object sizes make it less feasible to discern specific characteristics or object types, as illustrated in the Figure 30. Thus, having additional modalities and range estimation techniques such as optical triangulation [6] or passive radar [8] is needed to identify objects in many cases. Fortunately, for our toy outlier search, few small objects pass the high sinuosity criteria. During our manual examination we were able to associate mundane objects with the majority of trajectories with high sinuosity. The visual inspection of the recordings helped to group trajectories with high sinuosity into several categories: single birds (presumably hawks), flocks of birds, planes, leaves, and clouds. The supplementary material showcases typical reconstructed trajectories for all of these high sinuosity categories.

After manually scanning all trajectories flagged as outliers by the high sinuosity criteria, 144 trajectories remained ambiguous even after a manual check, with basic trajectory shapes similar to those generated in our synthetic dataset: looping segments, straight and piece-wise paths.

Finally, to establish the list of potential object categories present in the Dalek dataset, we manually labeled 12,000 trajectories and identified eight distinct categories shown in Figure 31. Among these, we define a category of "ambiguous" objects that were flying at a distance, making visual classification challenging. Despite the difficulty in categorizing these objects, they did not exhibit any abnormal behavior. Figure 32 also shows zoomed in infrared images of representative objects observed in this labeled dataset of trajectories: flocks of birds, the moon, planes, single birds, and clouds.




Figure 28. Left: Examples of reconstructed trajectory data points. Right: For each corresponding reconstructed object, we overlay three snapshots of the frame-by-frame object YOLO detections, inside green outlines and annotated with red text (unique identifiers). The snapshots are taken and located at the start, middle, and end of each trajectory. The background is a frame of the corresponding camera which does not contain any detected objects, overlaid for reference. Top row: objects with low reconstructed trajectory sinuosity, between 1.0 and 1.2. Middle row: objects with reconstructed trajectory sinuosity ranging from 1.2 to 3.0. Bottom row: objects with reconstructed trajectory sinuosity > 3.0 .




Figure 29. Left: Airplanes on Cam-7. Middle: Birds and clouds on Cam-4. Right: Helicopters, Moon on Cam-7. All are examples of detected objects that can be visually identified by inspecting the detection bounding box and the corresponding video recording. We show here examples of airplane, bird, cloud, helicopter and Moon detections. For each reconstructed trajectory, we overlay three snapshots of the frame-by-frame object YOLO detections, inside green outlines and annotated with red text (unique identifiers). The snapshots are taken and located at the start, middle, and end of each trajectory. The background is a frame of the corresponding camera which does not contain any detected objects.

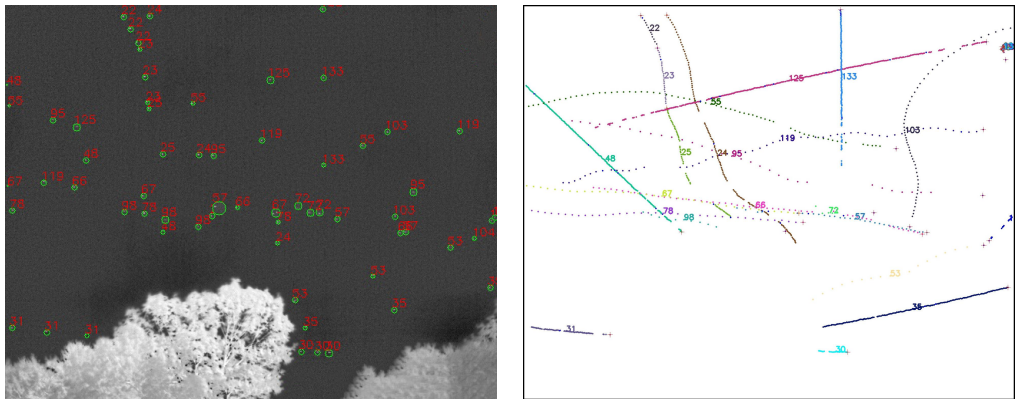


Figure 30. Examples of trajectories created by objects with small apparent size. Left: For each corresponding reconstructed trajectory, we overlay three snapshots of the frame-by-frame object YOLO detections, inside green outlines and annotated with red text (unique identifiers). The snapshots are taken at different points along the trajectories. Right: Corresponding examples of reconstructed trajectory data points.

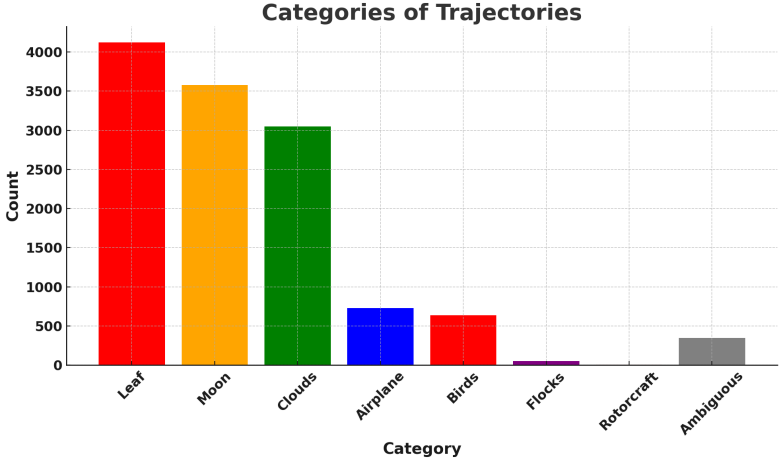


Figure 31. Manual classification of 12,000 reconstructed trajectories, a small subset of the full commissioning dataset, into 8 classes of objects.




Figure 32. After manual classification of reconstructed trajectories, we sample typical objects (in pairs) from each category. These images are crops of the objects for illustration purposes. First row: flocks of birds and the Moon; Second row: planes and single birds; Third row: clouds.

3.4.2. Likelihood analysis

Our goal in this section is to demonstrate the use of a likelihood-based test to characterize the statistical significance of the results of a quantitative outlier search. We will use as our test case the toy outlier search conducted in the previous section. This is a toy case because it is unimodal (one sensor type), degenerate (apparent 2D trajectory), and using only a single characteristic of the trajectory (“sinuosity”). In the future, we will have data from multiple modalities, use range estimation to remove the degeneracy (e.g. with triangulation), and introduce multiple characteristics from each modality (e.g. speed and size) using data from the full observatory. We follow the seminal paper [43] on likelihood-based tests of new physics.

In our case, we define the null hypothesis, H_0 , as describing a known dataset with no outliers. We test H_0 against the alternative hypothesis, H_1 , that our dataset has both this known background and outliers. In the toy outlier analysis described in the previous section, our observed dataset consists of measurements of apparent sinuosity of each reconstructed trajectory.

The p -value quantifies the statistical significance of the observations assuming the null hypothesis H_0 , i.e. the probability of observations of equal or worse compatibility with H_0 . A small p -value indicates that it would be unlikely for the null hypothesis alone to have generated the observed data.

Significance is defined in physics by a Z-score: $Z = \psi^{-1}(1 - p)$, where ψ^{-1} is the quantile of the standard Gaussian distribution of the values of the characteristic of interest, ϕ ; in other words, it is the inverse of its cumulative distribution.

The probability of observing a specific dataset consistent with a particular hypothesis is called the *likelihood*. In our toy case, ϕ , the physical parameter of interest for us, is trajectory apparent sinuosity. Note that when we go beyond our simple counting experiment here,

we will make a histogram \vec{n} of detected and reconstructed object characteristics, such as acceleration or trajectory inflexion points.

If we assume the signal s and background b have properties that are known perfectly, we have an ideal experiment and the likelihood is trivial:

$$\mathcal{L}(\vec{n}, \phi) = \prod_{\text{bins}} \text{Poisson}(n_i | \phi s_i + b_i)$$

In a realistic measurement however, we have to account for non-negligible statistical and systematic uncertainties. Sources of systematic uncertainties in our measurements can include for example: calibrations such as the conversion from 16-bit to temperature units scale, efficiencies in identification and reconstruction, resolution for example of the reconstructed acceleration. Let us consider for the sake of this simple example only one systematic uncertainty: assume we know the uncertainty of our efficiency for reconstructing trajectories of objects which can affect our selection, and assume that this uncertainty has a Gaussian distribution. We represent it with a parameter α of scale such that the values ± 1 correspond to this uncertainty. Then we have to modify the likelihood to model this uncertainty:

$$\mathcal{L}(\vec{n} | \phi, \alpha) = \prod_{\text{bins}} \text{Poisson}(n_i | \phi s_i + b_i(\alpha)) \cdot \text{Gauss}(0 | \alpha, 1)$$

where $b_i(\alpha)$ is a response function which tells us how much the background changes as a result of α , and $\text{Gauss}(0 | \alpha, 1)$ represents a Gaussian of mean 0 and variance 1 evaluated in α . In principle we can build an approximation of $b_i(\alpha)$ using the synthetic dataset for example, running our selection with $\pm 1\sigma$ variations of α . From now on we refer to this full likelihood as a *profile likelihood*.

$$\mathcal{L}(\vec{n} | \phi, \alpha) = \prod_{i=1}^N \frac{(\phi s_i + b_i(\alpha))^{n_i} e^{-(\phi s_i + b_i(\alpha))}}{n_i!} \cdot \text{Gauss}(0 | \alpha, 1)$$

Having defined our profile likelihood based on our measurement and its uncertainties, we can move onto a statistical test to compare the two hypothesis \mathcal{H}_0 and \mathcal{H}_1 . A classical frequentist approach is to use the likelihood-ratio test. For \mathcal{H}_0 the profile likelihood-ratio is defined as

$$\lambda(\phi) = \frac{\mathcal{L}(\vec{n} | \phi, \hat{\alpha})}{\mathcal{L}(\vec{n} | \hat{\phi}, \hat{\alpha})}$$

λ has values in $[0, 1]$. Being close to 1 reflects a good agreement between the measurement and the given value of ϕ . The numerator is a conditional maximal-likelihood estimator of α : $\hat{\alpha}$ is the value of α for which \mathcal{L} is maximal given a specific value of ϕ . $\hat{\phi}$ and $\hat{\alpha}$ are maximum-likelihood estimators. In practice, the test statistics used for setting an upper limit on ϕ are

$$q_\phi = \begin{cases} -2 \ln \lambda(\phi) & \hat{\phi} \leq \phi \\ 0 & \hat{\phi} > \phi \end{cases}$$

The higher q_ϕ , the less agreement there is between the measurement and the given value of ϕ . In order to find the sensitivity of our measurement for an exclusion limit, we have to look at the median significance with which we could reject $\phi > 0$, assuming a dataset generated under the null hypothesis of background only. The approximated median exclusion significance or Asimov discovery significance is, assuming the Wald approximation [43]:

$$\text{med}[Z_\phi | 0] = \sqrt{q_{\phi, A}}$$

where $q_{\phi, A} = -2 \ln \lambda_A(\phi)$ is the test statistics for the Asimov dataset, defined as the dataset where the bin counts are equal to their expectation values. If the corresponding p-value is below a threshold β , we say that the value of ϕ is excluded at a confidence level of $1 - \beta$. In practice the upper limit is found numerically as the value for which $p = \beta$.

We apply this method to the toy search for outliers described in the previous section, focusing on a simple counting experiment instead of a histogram experiment. We start by defining our background and signal counts, including uncertainties. Out of $N^{obs} = 502,015$ reconstructed trajectories, $N_{flag} = 81,873$ were flagged with a sinuosity greater than 3.0, and after manual inspection, $N_s^{obs} = 144$ trajectories remained ambiguous, although with plausible mundane explanations. Before we compute a profile likelihood, we have to estimate the uncertainties on these counts.

On the manually labeled real-world dataset containing a variety of objects, we found a recall for the object detection stage of $\sim 63\%$. An object detection recall of 63% means that we should assign an uncertainty of at least $\sim 37\%$ to the reconstructed sinuosity: we are potentially missing one out of three points in any given trajectory. The worst impact on the Euclidean distance from trajectory start to end happens when we are missing points at the start or at the end of a straight trajectory, i.e. a fractional uncertainty of 37%. For the trajectory curvilinear length estimation, the worst impact happens for high sinuosity trajectories: for example for a trajectory in the shape of saw teeth, the total length of the trajectory may be off by as much as 50%. Sinuosity is the ratio of the two, to which we assign a fractional uncertainty of 62% after propagating the two fractional uncertainties. In the toy outlier analysis, instead of selecting trajectories with sinuosity greater than 3 we should be using a threshold of 1.14 if we account for this uncertainty. Thus, the number of flagged trajectories is assigned an uncertainty of $\sim 400\%$. The manual inspection identifies a fraction $\alpha = \frac{N_s^{obs}}{N_{flag}}$ as ambiguous. Since manual processes are not flawless, we assign α an uncertainty of 10%. When propagating errors we use $\Delta N_s^{obs} = N_{flag} \Delta \alpha + \alpha \Delta N_{flag}$. We find that the signal count is $N_s^{obs} = 144 \pm 8763$.

On a synthetic dataset without any clouds, rain or high relative humidity, the identification recall metric for SORT was estimated at 87% in Section 2.3.3. To be conservative, we set the overall efficiency of the pipeline (the fraction of true objects within our acceptance that are reconstructed), at $\epsilon = 80 \pm 20\%$. After propagating the uncertainty from ϵ we unfold the number of background trajectories to $N_b = \frac{N^{obs} - N_s^{obs}}{\epsilon} = 627,339 \pm 324,668$ and the number of outlier trajectories to $N_s = \frac{N_s^{obs}}{\epsilon} = 180 \pm 10,999$. This is our estimate of true background and the count of outlier objects (as defined by our toy outlier criteria) in the parameter space of apparent 2D trajectory sinuosity, within the Dalek's acceptance, if the detection pipeline was perfect. Finally we use the Python package `hepstats` [44] to implement a likelihood inference using the Asimov dataset. We assume that the uncertainties on N_s and N_b have a Gaussian shape. We find an observed upper limit of 18,271 at 95% confidence level for the number of outlier trajectories as shown in Figure 33.

4. Discussion

The $\sim 500,000$ reconstructed aerial trajectories that we report from this commissioning period of five months give us insight into the volume of object parameter space that the Dalek is sensitive to. From the commissioning studies, we expect that the Dalek acceptance includes objects of apparent size larger than 3 pixels, which corresponds to a standard airliner of wingspan 50 m at a distance of 10 km or a bird of wingspan 1.5 m at a distance of 300 m. Kinematics-wise, objects with a speed larger than 200 m/s at a distance of 1 km or a speed larger than 2 km/s at a distance of 10 km are outside the Dalek's acceptance based on both the frame rate and our analysis requirement of at least four points in a trajectory. The reconstructed aerial trajectories are analyzed with a toy outlier selection based on high sinuosity. About 16% of trajectories are flagged as outliers and manually examined. 144 ambiguous objects remained, which are likely mundane objects but cannot be elucidated without distance estimation and multi-modality at this stage of development. We demonstrate the application of a likelihood-based statistical test to evaluate the significance of this toy outlier analysis. Our observed count of ambiguous outliers combined with systematic uncertainties yields an upper limit of 18,271 outliers (toy outliers flagged as unusual) having passed through the Dalek's detection volume and




Figure 33. Graphical representation of the observed and expected upper limit calculation. This shows the upper limit on the ambiguous outlier count at 0.05 significance level (shown as a red line), for a counting experiment where $N_b = 627,339 \pm 324,668$ and $N_s = 180 \pm 10,999$, assuming gaussian uncertainties. The observed upper limit at 95% confidence level is 18,271.

time window during the commissioning period at a 95% confidence level. This upper limit corresponds to $\sim 4\%$ of the objects we reconstructed and $\sim 22\%$ of the objects flagged as outliers based on their apparent trajectory sinuosity. For context, classified studies done by governmental agencies such as the All-domain Anomaly Resolution Office (AARO), which are often able to leverage distance estimations from radar data and data from multiple modalities, reported that $\sim 3\%$ of the cases brought to their attention remained ambiguous. We expect that our upper limit will decrease in the future after we improve our detection pipelines, and include multiple instruments and range estimation in the survey. This paper also paves the way for conducting a similar commissioning process for each of our observatory's instruments before conducting a in-depth, multi-modal aerial survey. We plan to apply this likelihood-based method to all of our future outlier searches.

Author Contributions: Conceptualization, A.D., L.D., S.L., W.W.; methodology, Er.K. and M.S.; software, E.M., R.C., A.F., M.S. and W.W.; formal analysis, L.D., R.C., A.F. and A.B.; resources, A.D., F.S. and A.W.; data curation, L.J.; writing—original draft preparation, L.D., M.S., R.C., A.F., L.J., S.L., and F.S.; writing—review and editing, M.P., M.S., S.L., and W.W.; supervision, A.L., W.W. and S.L.; project administration, Ez.K. and A.L.; funding acquisition, A.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by private donations to the Galileo Project.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: The authors thank Cameron Pratt for his help with review and editing.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Office of the Director of National Intelligence. 2022 Annual Report on Unidentified Aerial Phenomena. <https://www.dni.gov/index.php/newsroom/reports-publications-2023/3667-2022-annual-report-on-unidentified-aerial-phenomena>, 2022.
2. Office of the Director of National Intelligence. 2023 Consolidated Annual Report on Unidentified Anomalous Phenomena. <https://www.dni.gov/index.php/newsroom/reports-publications-2023/3733-2023-consolidated-annual-report-on-unidentified-anomalous-phenomena>, 2023.
3. Team, N.U.A.P.I.S. NASA Unidentified Anomalous Phenomena Independent Study Team Report. Final report, NASA, 2023.
4. Watters, W.A.; Loeb, A.; Laukien, F.; Cloete, R.; Delacroix, A.; Dobroshinsky, S.; Horvath, B.; Kelderman, E.; Little, S.; Masson, E.; et al. The Scientific Investigation of Unidentified Aerial Phenomena (UAP) using multimodal ground-based observatories. *Journal of Astronomical Instrumentation* **2023**, *12*, 2340006.
5. Mead, A.; Little, S.; Sail, P.; Tu, M.; Watters, W.A.; White, A.; Cloete, R. Multi-band acoustic monitoring of aerial signatures. *Journal of Astronomical Instrumentation* **2023**, *12*, 2340005.
6. Szenher, M.; Delacroix, A.; Keto, E.; Little, S.; Randall, M.; Watters, W.A.; Masson, E.; Cloete, R. A hardware and software platform for aerial object localization. *Journal of Astronomical Instrumentation* **2023**, *12*, 2340002.
7. Cloete, R.; Bridgham, P.; Dobroshinsky, S.; Ezell, C.; Fedorenko, A.; Laukien, F.; Little, S.; Loeb, A.; Masson, E.; Szenher, M.; et al. Integrated Computing Platform for Detection and Tracking of Unidentified Aerial Phenomena (UAP). *Journal of Astronomical Instrumentation* **2023**, *12*.
8. Randall, M.; Delacroix, A.; Ezell, C.; Kelderman, E.; Little, S.; Loeb, A.; Masson, E.; Watters, W.A.; Cloete, R.; White, A. Skywatch: A passive multistatic radar network for the measurement of object position and velocity. *Journal of Astronomical Instrumentation* **2023**, *12*, 2340004.
9. Ailleris, P. Exploring Unidentified Aerospace Phenomena through Instrumented Field Studies: Historical Insights, Current Challenges, and Future Directions. *Limina-The Journal of UAP Studies* **2024**, *1*, 11–30.
10. Vodopivec, A.; Kayal, H. Autonomous sensor network for detection and observation of meteors. In Proceedings of the Proceedings of the International Meteor Conference, 2018, p. 156.
11. Skycam5. <https://www.informatik.uni-wuerzburg.de/aerospaceinfo/wissenschaft-forschung/skycam-5/>, 2020.
12. Kayal, H. HYPER-SETI—A New Way of Searching for Extraterrestrial Intelligence. In Proceedings of the 70th International Astronautical Congress (IAC), Washington DC, United States, 2019, pp. 21–25.
13. Szydagis, M.; Knuth, K.; Kugielsky, B.; Levy, C.; McGowan, J.; Phelan, M.; Voorhis Jr, G. Initial Results From the First Field Expedition of UAPx to Study Unidentified Anomalous Phenomena. *arXiv preprint arXiv:2312.00558* **2023**.
14. Colas, F.; Zanda, B.; Bouley, S.; Jeanne, S.; Maltgoire, A.; Birlan, M.; Blanpain, C.; Gattacceca, J.; Jorda, L.; Lecubin, J.; et al. FRIPON: a worldwide network to track incoming meteoroids. *Astronomy & Astrophysics* **2020**, *644*, A53.
15. Sigma2 progress report 2021. <https://www.3af.fr/agenda/download-the-sigma2-progress-report2149>, 2021.
16. Villarroel, B.; Marcy, G.W.; Geier, S.; Streblyanska, A.; Solano, E.; Andruk, V.N.; Shultz, M.E.; Gupta, A.C.; Mattsson, L. Exploring nine simultaneously occurring transients on April 12th 1950. *Scientific reports* **2021**, *11*, 12794.
17. Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A.C.; Lo, W.Y.; et al. Segment anything. In Proceedings of the Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 4015–4026.
18. Gstreamer: Open source multimedia framework. <https://gstreamer.freedesktop.org/>, 2022.
19. Zhang, Z. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2000**, *22*, 1330–1334. <https://doi.org/10.1109/34.888718>.
20. Bradski, G. The OpenCV Library. *Dr. Dobb's Journal of Software Tools* **2000**.
21. Boson image non-uniformities (INUs): Application note. <https://groupgets-les.s3.amazonaws.com/boson/documents/Boson>
22. Boson lens calibration: Application note. <https://pdfslide.net/documents/application-note-amazon-web-services-lens-calibration-boson-lens-calibration-application.html>, 2017.
23. Van Diggelen, F.; Enge, P. The world's first GPS MOOC and worldwide Laboratory using Smartphones. In Proceedings of the Proceedings of the 28th international technical meeting of the satellite division of the institute of navigation (ION GNSS+ 2015), 2015, pp. 361–369.
24. Federal Aviation Administration. *Automatic Dependent Surveillance-Broadcast (ADS-B) Out equipment and use.*, 2024.
25. Federal Aviation Regulations, Part 91, 91.227. .
26. Schäfer, M.; Strohmeier, M.; Lenders, V.; Martinovic, I.; Wilhelm, M. Bringing up OpenSky: A large-scale ADS-B sensor network for research. In Proceedings of the IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks. IEEE, 2014, pp. 83–94.
27. Marchand, E.; Uchiyama, H.; Spindler, F. Pose Estimation for Augmented Reality: A Hands-On Survey. *IEEE Transactions on Visualization and Computer Graphics* **2016**, *22*, 2633 – 2651. <https://doi.org/10.1109/TVCG.2015.2513408>.
28. Jocher, G.; Chaurasia, A.; Stoken, A.; Borovec, J.; Kwon, Y.; Michael, K.; Fang, J.; Wong, C.; Yifu, Z.; Montes, D.; et al. ultralytics/yolov5: v6. 2-yolov5 classification models, apple m1, reproducibility, clearml and deci. ai integrations. *Zenodo* **2022**.
29. Jiang, P.; Ergu, D.; Liu, F.; Cai, Y.; Ma, B. A Review of Yolo algorithm developments. *Procedia computer science* **2022**, *199*, 1066–1073.
30. Liu, D.C.; Nocedal, J. On the limited memory BFGS method for large scale optimization. *Mathematical programming* **1989**, *45*, 503–528.

31. Passman, S.; Larmore, L. Atmospheric Transmission. *Rand Paper* **1956**, p. 897.
32. Minkina, W.; Klecha, D. Atmospheric transmission coefficient modelling in the infrared for thermovision measurements. *Journal of sensors and sensor systems* **2016**, *5*, 17–23.
33. Gaussorgues, G. *Infrared Thermography*; Springer Science+Business Media, B. V., Dordrecht, 1994; p. 552.
34. Minkina, W.; Dudzik, S. *Infrared thermography: errors and uncertainties*; John Wiley & Sons, 2009.
35. Bewley, A.; Ge, Z.; Ott, L.; Ramos, F.; Upcroft, B. Simple online and realtime tracking. In Proceedings of the 2016 IEEE International Conference on Image Processing (ICIP), 2016, pp. 3464–3468. <https://doi.org/10.1109/ICIP.2016.7533003>.
36. Verbraak, T.; Ellerbroek, J.; Sun, J.; Hoekstra, J. Large-scale ADS-B data and signal quality analysis. In Proceedings of the 12th USA/Europe Air Traffic Management Research and Development Seminar, 2017.
37. Ali, B.S.; Schuster, W.; Ochieng, W.Y. Evaluation of the capability of automatic dependent surveillance broadcast to meet the requirements of future airborne surveillance applications. *The Journal of Navigation* **2017**, *70*, 49–66.
38. Open-Meteo API. <https://open-meteo.com/>.
39. Bernardin, K.; Stiefelhagen, R. Evaluating multiple object tracking performance: the CLEAR MOT metrics. *URASIP Journal on Image and Video Processing* **2008**, *12*, 2340006.
40. Milan, A.; Leal-Taixe, L.; Reid, I.; Roth, S.; Schindler, K. Mot16: A benchmark for multi-object tracking. *URASIP Journal on Image and Video Processing* **2016**, *1*.
41. Nevatia, Y.L.C.H.R. Learning to associate: HybridBoosted multi-target tracker for crowded scene. *IEEE Conference* **2009**.
42. Ristani, E.; Solera, F.; Zou, R.S.; Cucchiara, R.; Tomasi, C. Performance Measures and a Data Set for Multi-Target, Multi-Camera Tracking. *IEEE Conference* **2016**, *2*.
43. Cowan, G.; Cranmer, K.; Gross, E.; Vitells, O. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C* **2011**, *71*, 1–19.
44. hepstats. <https://scikit-hep.org/hepstats/>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Hours	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
Camera 1																								
Summer																								
Winter																								
Equinox																								
Recording																								
Camera 2																								
Camera 3																								
Summer																								
Winter																								
Equinox																								
Recording																								
Camera 4																								
Summer																								
Winter																								
Equinox																								
Recording																								
Camera 5																								
Summer																								
Winter																								
Equinox																								
Recording																								
Camera 6																								
Summer																								
Winter																								
Equinox																								
Recording																								
Camera 7																								
Summer																								
Winter																								
Equinox																								
Recording																								
Camera 8																								
Summer																								
Winter																								
Equinox																								
Recording																								

Table A1. Hours of the day when the sunshades are mechanically closed in summer (yellow - May, June, July), winter (blue - November, December, January) and equinox (green - all other months) modes, in UTC-4h timezone. In orange, the schedule shows the hours when we are *not* recording, in local timezone which is UTC-4h in the summer and UTC-5h in the winter due to daylight savings. In future we will align both schedules to be identical, but during this commissioning period of five months each camera had a different amount of "lost" recording time, i.e. recordings of blank frames because of closed sunshades.

Supplementary Materials: Commissioning An All-Sky Infrared Camera Array for Detection Of Airborne Objects

Laura Domine^{1,3,*}, Ankit Biswas³, Richard Cloete^{1,3}, Alex Delacroix^{1,3}, Andriy Fedorenko³, Lucas Jacaruso³, Ezra Kelderman³, Eric Keto^{1,3}, Sarah Little^{2,3,4}, Abraham Loeb^{1,3,†}, Eric Masson³, Mike Prior³, Forrest Schultz^{3,5}, Matthew Szenher³, Wes Watters^{2,3}, Abby White^{1,3}

1. Examples of trajectories manually identified

We hand picked some representative trajectories from our commissioning dataset that we manually identified as birds, leaves or clouds.

1.1. Birds

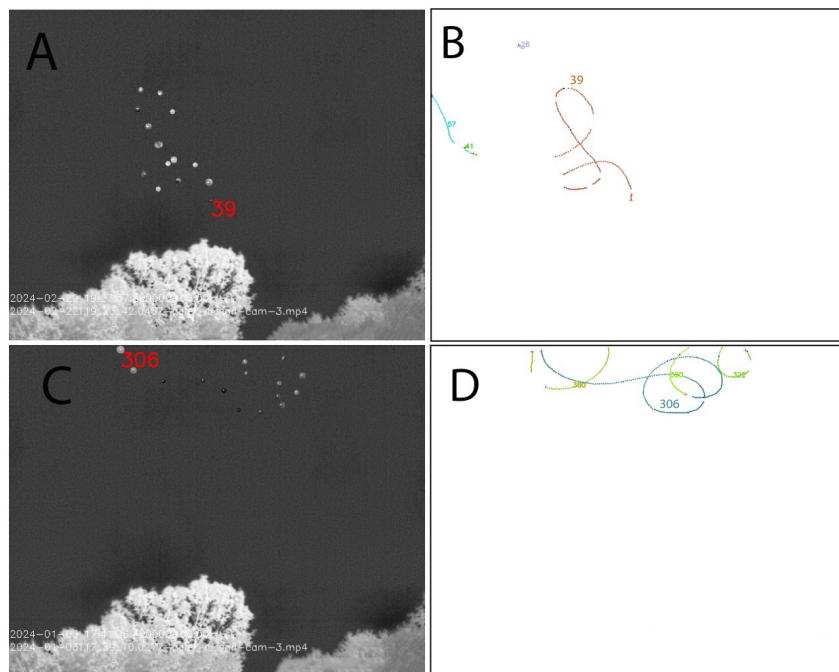


Figure S1. Left: Camera background image, overlaid with snapshots from the objects' detections along their trajectories, for two specific trajectories with ID 39 and 306, identified as hawks after manual examination. Right: All reconstructed trajectories in a given day, including trajectories 39 and 306, summarized by assigning a unique color and identifier to each point of these trajectories.

Figure S1 shows trajectories with IDs 39 and 306 that include multiple loops (sinuosity > 3). The object exhibits an unusually long trajectory lifespan. The trajectory with ID 306 re-entered the scene and exhibited a similar behavior, generating additional looping trajectories such as those labeled 322, 320, and 300. Visual inspection identifies these objects as birds, presumably hawks or eagles that loop around thermals and may hover for a short period of time. The next Figure S2 presents another looping trajectory (Figure S2C and Figure S2D). The same Figure also showcases an interesting object with sinuosity < 3 in Figure S2A and Figure S2B: the trajectory with ID 14 starts in the upper right corner (ID 12) and then moves sharply downward, decelerating as it transitions into a straight trajectory labeled ID 14. This motion suggests a change in velocity, marked by an initial curved path followed by a more linear motion pattern.

Figure S3 presents another example of a looping trajectory. The object makes a sharp U-turn and moves backward. We also identify objects traveling in pairs or as part of a

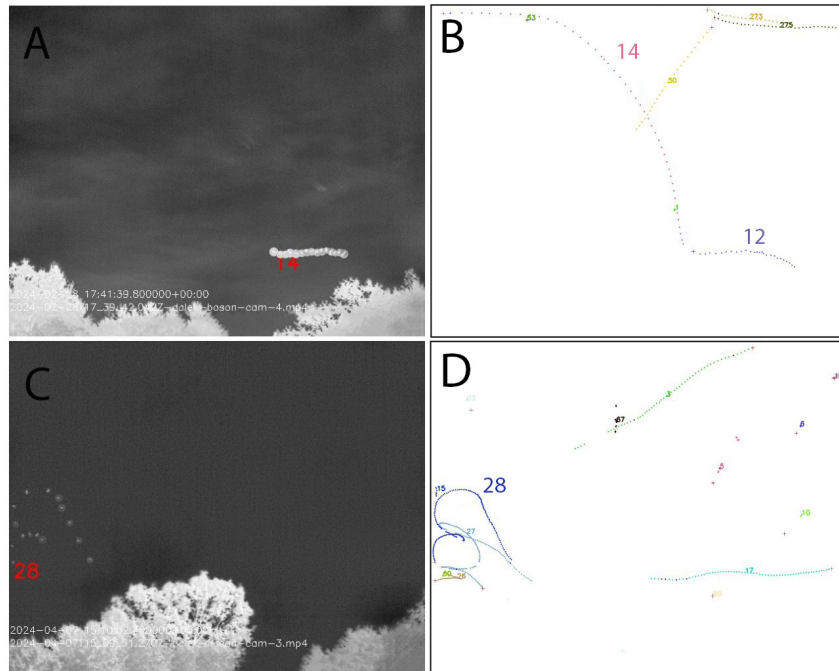


Figure S2. Left: Camera background image, overlaid with snapshots from the objects' detections along their trajectories, for two specific trajectories with ID 14 and 28, identified as hawks after manual examination. Right: All reconstructed trajectories in a given day, including trajectories 14 and 28, summarized by assigning a unique color and identifier to each point of these trajectories.

larger formation, changing their alignment as they move. These are identified as flocks of birds, illustrated in Figure S4.

1.2. Leaves

We also find interesting trajectories where the objects change their shape from frame to frame, such as on Figure S5. Visual inspection of the recordings helps to identify trajectories 37, 115, and 18 as leaves flying off and toward the camera.

1.3. Clouds

The next set of objects are clouds of different shapes. Manual labeling shows that 35% of our trajectories are made by clouds, which are shown in Figure S6.

2. Order of magnitude range estimation using common known object speeds

All-sky cameras alone face several limitations due to the lack of range data, which affects our ability to accurately interpret and detect objects and their trajectories. The scale ambiguity makes it difficult to differentiate between nearby and distant objects (such as those in Figure S4), leading to challenges in estimating their size, speed, or altitude. Distant objects may appear blurred, and atmospheric distortions further degrade image clarity, making detection harder. Additionally, without range information, it is a challenge to accurately track object speed or compensate for motion blur, especially for apparently fast-moving objects like bugs or birds flying very close to the camera. For example, in Figure S5B trajectories with IDs 18, 37, and 15 show disrupted, looping shapes. In this case, visual inspection resolves the nature of the object (leaves), but even this simple approach fails when we need to differentiate objects measuring only a few pixels across. However, given an observed angular speed and an estimate of the most common average speed of given objects classes, we can approximately estimate the order of magnitude for the range of the object from the camera. For example, in Figure S5A, C, and E, leaves cross a 50-degree field of view in 15 seconds, 5 seconds, and 3 seconds, respectively, with wind-driven speeds typically ranging from 2–5 m/s. By combining the observed crossing times and speeds,




Figure S3. Left: Camera background image, overlaid with snapshots from the objects' detections along their trajectories, for the specific trajectory with ID 812, identified as a hawk after manual examination. Right: All reconstructed trajectories in a given day, including trajectory 812, summarized by assigning a unique color and identifier to each point in these trajectories.




Figure S4. Zoomed view of the individual object detections overlaid along different reconstructed trajectories. The red number is the unique identifier assigned to each trajectory by the tracking algorithm (SORT). All of these examples are flocks of birds after manual examination.

we estimate the approximate order of magnitude for distances from the camera to the leaves, which are 5, 11, and 35 meters, respectively. Similarly, taking the average speed of birds (e.g., hawks when flying) as 10–16 m/s, the time needed to cross the field of view in Figure S1A and C, which is 10 and 18 seconds, gives distances of approximately 150 and 270 meters, respectively. In Figure S4, flocks of bird corresponding to the trajectories with ID 246, 86, and 20, are estimated to be at a range of about 1.3 km, crossing the field of view within 42 seconds with the speed of flock of geese to be about 60 mph (25 m/s). To provide a more accurate estimate of the distance to the object, range estimation techniques such as optical triangulation [1] or passive radar [2] would be required. Both are in development for our observatory.

References

1. Szenher, M.; Delacroix, A.; Keto, E.; Little, S.; Randall, M.; Watters, W.A.; Masson, E.; Cloete, R. A hardware and software platform for aerial object localization. *Journal of Astronomical Instrumentation* **2023**, *12*, 2340002.
2. Randall, M.; Delacroix, A.; Ezell, C.; Kelderman, E.; Little, S.; Loeb, A.; Masson, E.; Watters, W.A.; Cloete, R.; White, A. Skywatch: A passive multistatic radar network for the measurement of object position and velocity. *Journal of Astronomical Instrumentation* **2023**, *12*, 2340004.

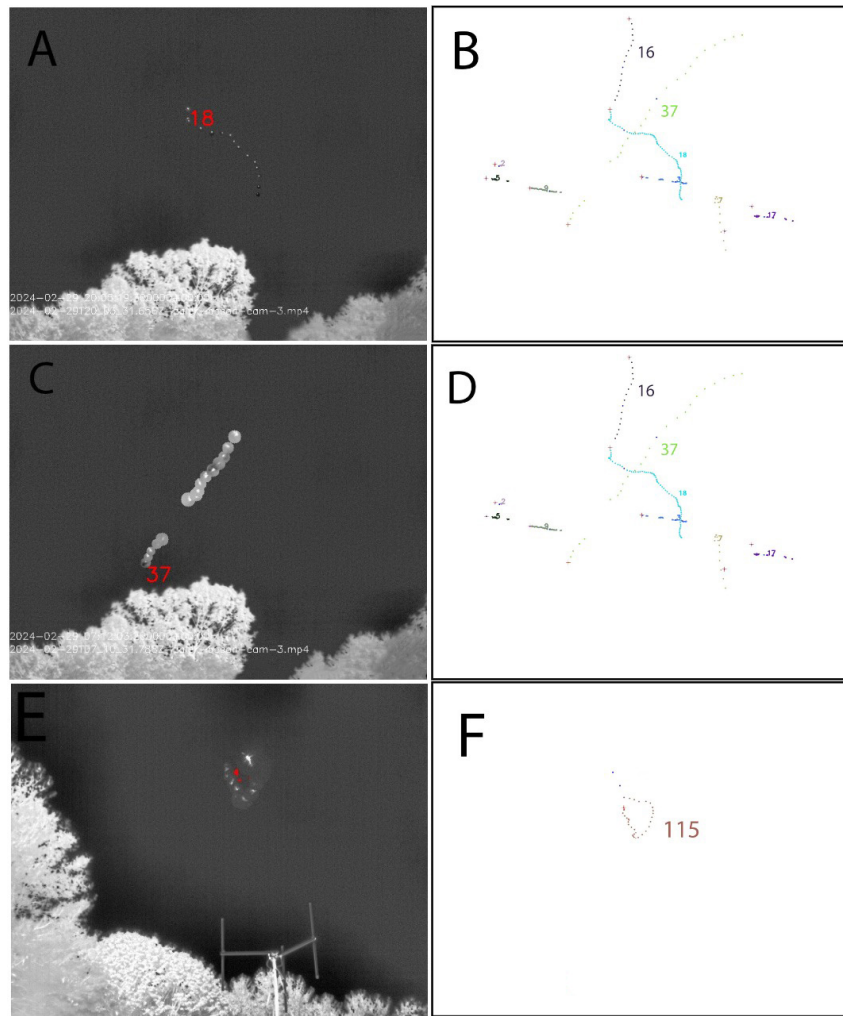


Figure S5. Left: Camera background image, overlaid with snapshots from the objects' detections along their trajectories, for three specific trajectories with ID 18, 37 and 115, which were all identified after manual examination as leaves. Right: All reconstructed trajectories in a given day, including trajectories 18, 37 and 115, summarized by assigning a unique color and identifier to each point of these trajectories.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.




Figure S6. Zoom onto the individual object detections overlaid along different reconstructed trajectories. The red number is the unique identifier assigned to each trajectory by the tracking algorithm (SORT). All of these examples are clouds after manual examination.