



Successful Recovery of an Observed Meteorite Fall Using Drones and Machine Learning

Seamus L. Anderson , Martin C. Towner , John Fairweather, Philip A. Bland, Hadrien A. R. Devillepoix , Eleanor K. Sansom , Martin Cupák, Patrick M. Shober , and Gretchen K. Benedix
Space Science and Technology Centre, School of Earth and Planetary Science, Curtin University, 314 Wark Avenue, Bentley, WA 6102, Australia
seamus.anderson@postgrad.curtin.edu.au

Received 2022 February 16; revised 2022 March 12; accepted 2022 March 17; published 2022 May 27

Abstract

We report the first-time recovery of a fresh meteorite fall using a drone and a machine-learning algorithm. The fireball was observed on 2021 April 1 over Western Australia by the Desert Fireball Network, for which a fall area was calculated for the predicted surviving mass. A search team arrived on-site and surveyed 5.1 km² area over a 4 day period. A convolutional neural network, trained on previously recovered meteorites with fusion crusts, processed the images on our field computer after each flight. Meteorite candidates identified by the algorithm were sorted by team members using two user interfaces to eliminate false positives. Surviving candidates were revisited with a smaller drone, and imaged in higher resolution, before being eliminated or finally being visited in person. The 70 g meteorite was recovered within 50 m of the calculated fall line, demonstrating the effectiveness of this methodology, which will facilitate the efficient collection of many more observed meteorite falls.

Unified Astronomy Thesaurus concepts: [Meteorites \(1038\)](#); [Fireballs \(538\)](#); [Convolutional neural networks \(1938\)](#)

1. Introduction

Besides recording the conditions of the protoplanetary nebula and the early solar system, meteorites also offer insights into the contemporary physical and chemical compositions of asteroids and other terrestrial bodies (Cuzzi et al. 2008; Nakamura et al. 2011). Some of these meteorites fall in regions on Earth where fireball observatory networks are active, making it possible to record the trajectory of the fireball as it ablates material from the originating meteoroid. For some fireballs, these data can then be used to simulate both forward and backward in time to predict where the resulting meteorite landed on Earth and where the meteoroid originated in the solar system. Thus, recovering and analyzing these “orbital meteorites” with constrained, prior orbits provides an incredibly unique insight into the geology and dynamic behavior of the asteroid belt and the nature of mass transfer between the belt and the inner solar system. The Desert Fireball Network (DFN; Bland et al. 2012; Howie 2017) is one of many research groups (Oberst et al. 1998; Olech et al. 2006; Spurný et al. 2006; Trigo-Rodríguez et al. 2006; Colas et al. 2015) within the Global Fireball Observatory (Devillepoix et al. 2020) that makes this possible.

In the past, recovering meteorites within a predicted area often consisted of four to six people walking 5–10 m apart, sweeping the area until the meteorite is found. This was labor intensive and suffered from a relatively low success rate of ~20%, considering how many of our trips (~40) have returned with one of our eight meteorites. This was partially due to the fact that an entire fall zone could rarely be covered in one trip. This has prompted efforts by multiple groups to recover meteorites using drones and machine learning, which massively reduces the time and labor required (Citron et al. 2017; Zender et al. 2018; AlOwais et al. 2019; Citron et al. 2021). Our previous efforts (Anderson et al. 2019, 2020) have shown

the ability of our methodology to positively identify already-recovered fresh meteorites, awaiting an opportunity to fully test our searching strategy.

On the night of 2021 April 1 such an opportunity presented itself as a meteorite fell over the Western Nullarbor on the Lintos Paddock of Kybo Station, Western Australia (Figure 1). Unfortunately, only two DFN camera stations, both to the east, were able to capture the event that caused the relatively short fall line to have considerable longitudinal uncertainty, expanding the area with a 90% likelihood of containing the meteorite to a total of 5.1 km² (Figure 1), with a predicted mass of the meteorite being between 150 and 700 g. These fall conditions were promising enough to warrant a field trip to survey the entire fall zone with a drone. The first 3 days we spent on-site consisted of surveying with a drone, and processing data with our machine-learning algorithm. On the fourth and final day we visited meteorite candidates with the drone and in person, and recovered the 70 g meteorite.

2. Methods

2.1. Fireball Observations and Modeling

Although the detailed analysis of the fireball and its orbit will be the subject of a future publication, here we briefly explain the data and methods that allowed us to predict the searching area. These analyses were done prior to the recovery, and have not been revised since. On 2021 April 1, two Desert Fireball Network observatories imaged a bright 3.1 s fireball (Figure 1), which was reported to the DFN team by our automated detection software (Towner et al. 2020), while astrometric calibration is performed following the methodology detailed in Devillepoix et al. (2018).

In total, 78 data points were recorded from just two observatories located at Mundrabilla station and O’Malley siding, 149 km and 471 km from the end point, respectively (Figure 1). The nominal trajectory started at 87 km altitude at 25.4 km s⁻¹, and the bolide was observed down to 25 km at 8.4 km s⁻¹, on a 64° slope. Because of the distance of the viewpoints and a low convergence angle of planes of 28°,

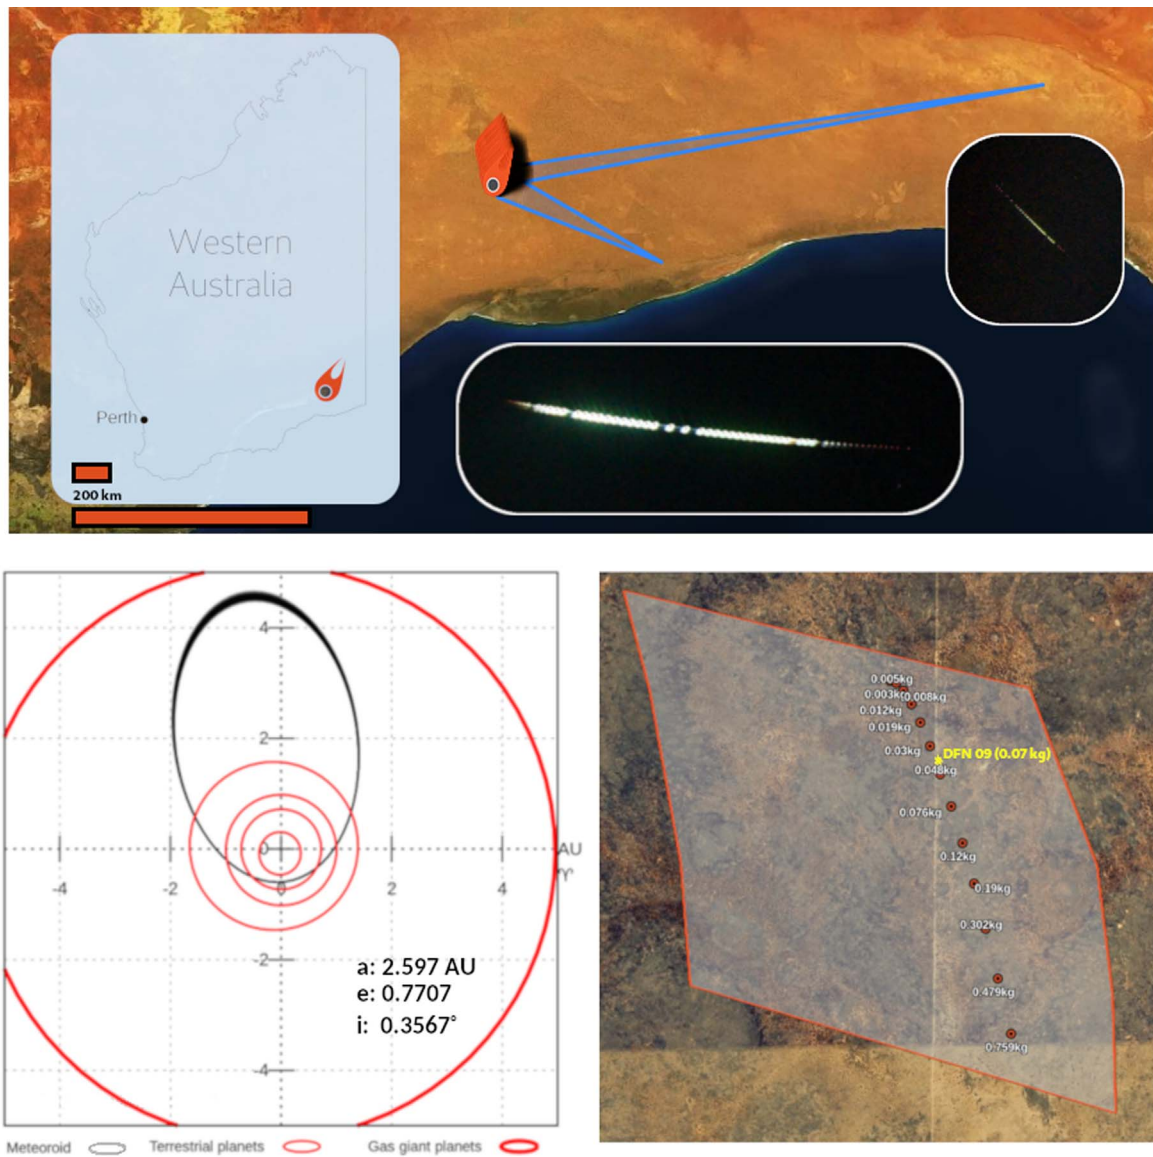


Figure 1. The DFN 09 meteorite fall at Kybo Station, Western Australia. Clockwise from top: Fireball observations from DFN camera stations at Mundrabilla Station and O’Malley Siding, and their location within WA. The 5 km², 90% certainty searching area (transparent white), the best-fit fall line (red markers), and the location of the recovered meteorite (yellow star). Preimpact orbit for the DFN 09 meteoroid ($a = 2.597 \pm 0.017$ au, $e = 0.7707 \pm 0.0018$, $i = 0^\circ 3567 \pm 0^\circ 0053$, $\omega = 85^\circ 862 \pm 0^\circ 015$, and $\Omega = 191.90282 \pm 0^\circ 00073$ 39 (1σ uncertainties)).

observation conditions were not ideal, resulting in a poorly constrained trajectory. To estimate the variability in the trajectory, in a similar manner to Devillepoix et al. (2022), we use a Monte Carlo approach, randomizing astrometric observations within errors, generating 1000 clones. This analysis shows a 500 m standard deviation on the end point.

For estimating the surviving mass of the object, we use the alpha–beta method from Gritsevich et al. (2012) and Sansom et al. (2019). The noisy velocity data resulted in a poorly constrained surviving mass between 150 and 700 g, assuming a spherical to rounded brick shape, a bulk density of 3.5 g cm^{−3}, and a shape change parameter of 2/3. We model the atmospheric conditions numerically using the Weather Research and Forecasting (WRF) model version 4.0 with the Advanced Research WRF dynamic solver (Skamarock et al. 2021) with three model runs starting on 2021 April 1 at 00UT, 06UT, and 12UT, all giving similar models.

Using these atmosphere models, we then propagate bright flight observations to the ground using the dark flight model from Towner et al. (2022). The uncertainty on the mass, and the uncertainty on the positions, are the dominant sources of uncertainty. Thanks to the dominant wind being more or less in the same direction as the fireball, the small mass end of the fall line was somewhat compacted. From these simulations, we defined two search areas at different confidence levels: a 90% confidence level yielding a search area of 5 km² (Figure 1), and a 99% confidence level that gives a significantly larger search area of 8 km².

To compute the pre-encounter orbit, we start by modeling the initial speed with an Extended Kalman Filter/Smother (Sansom et al. 2015), applied to the nominal trajectory. Using the integrator from Jansen-Sturgeon et al. (2019), we then propagate the position of the meteoroid backward until it is 10× outside the sphere of influence of the Earth–Moon system.

The position is then propagated forward to the date of impact, ignoring the influence of the Earth and the Moon. From this point we convert position/velocity to ecliptic orbital elements (J2000). To quantify errors, we repeat this process using Monte Carlo randomization of the initial velocity vector within uncertainties. The resultant preimpact, orbital parameters for this meteoroid were (Figure 1) semimajor axis $a = 2.597 \pm 0.017$ au, eccentricity $e = 0.7707 \pm 0.0018$, inclination $i = 0^\circ.3567 \pm 0^\circ.0053$, argument of perihelion $\omega = 85^\circ.862 \pm 0^\circ.015$, and longitude of ascending node $\Omega = 191^\circ.90282 \pm 0^\circ.00073$ 39 (1σ uncertainties).

2.2. Drone Surveying and Machine Learning

Our field-based work is a continuation of the methodology presented in Anderson et al. (2020). For this trip we used a DJI M300 drone with a Zenmuse P1 camera (44 MP) to survey the 5.1 km² fall line at 1.8 mm pixel⁻¹ with 20% overlap among images in each direction, which took between 2.5 and 3 days to complete. This ground sampling distance would image the meteorite at an apparent size between 15 and 65 pixels in diameter, due to inherent uncertainties in meteoroid properties during the fall. If the meteorite were predicted to have a significantly larger mass, we would increase the survey height to keep the same apparent size. We processed the data on-site using a desktop computer with an RTX 2080 Ti GPU. Our algorithm could process one flight’s worth of images (~30 minutes) in approximately 65 minutes.

Our algorithm works by taking a full 44 MP image and splitting it into tiles (125 × 125 pixels) with a 70 pixel overlap in each direction, to ensure the meteorite has a chance to appear fully in at least one tile. Each tile is fed into the binary image classifier presented in Anderson et al. (2020), constructed using python and keras (Chollet et al. 2015), which scores each tile from 0 (nonmeteorite) to 1 (meteorite). To obtain True (meteorite) training data, we relied on our library of meteorite images from previous trips as well as taking new images of the real meteorites: Camel Donga (Eucrite), Mulga North (H6), and Wiluna (H5) (supplied by the Western Australian Museum), at the fall zone. In total, we used 28 individual stones from these three meteorite falls, ranging in size from 2 to 14 cm (on their longest dimension). We repeated meteorite image collection each day and for each weather condition that affected illumination. Our True pool was comprised of all the meteorites we imaged on-site and an equal number again sourced from our library, totaling ~100,000 tiles. To create False tiles (non-meteorites) we took images from the surveyed fall zone at random, checked to ensure there were no meteorites in the frame, and split them into tiles. These totaled to more than 1 M depending on the day of the flights being sampled.

To maintain a balanced training set, while also sampling as much of our data set as possible, we employed what we call “rotation training.” To form a training set, we use a constant 80% of our True pool (with 20% kept for validation) and randomly select an equal number from the False pool, then we train for five epochs (rounds of training). After this we deselect these False tiles and randomly select a new False set from the pool, then repeat the training process until the model had worked through the False pool twice. The validation set contained only True tiles, so that we could monitor the most important performance metric: meteorite detection chance. We initially trained the model until it rotated through the False pool twice. By the end of training, we achieved a training accuracy

of 99.93%, and a validation accuracy (meteorite detection chance) of 91%.

As we predicted on each flight from our first day of surveying, we monitored the quality of the prediction by viewing the distribution of confidence values for a given image, in the form of a histogram (Figure 2). When the distribution for an image resembled that of Figure 2(A), we were satisfied, while images that produced a histogram like Figure 2(B) were flagged as “problem images.” We believe they contain physical features from the survey area that were not yet included in the training data, causing the model to infer generously, creating more unnecessary false positives to sort later in the process. We inspected some of these images, split them into tiles, and added them to the False pool. We later retrained on this augmented False pool, though only for two epochs and one rotation. Using the retrained model we would repredict on these problem images, which usually resulted in a more palatable confidence distribution.

We inspected meteorite candidate tiles in four stages (Figure 3). The first stages were identified by the model to have a confidence >0.7, and we inspected them using the 3 × 3 grid graphical user interface (GUI) described in Anderson et al. (2020), which was ideal for eliminating obvious false positives. The user would be given a 3 × 3 grid displaying nine tiles, with some being sourced from the True pool as tests for the user, to characterize their own performance. They identified interesting tiles by typing the corresponding key on the number pad before moving to the next set. Uninteresting tiles were added to the False pool, while those of interest became second stages that were further inspected in a separate image viewer that allowed pan and zoom control over the whole image. If the candidate was still of interest, it became a third stage. To inspect the third stages, we compiled all of their GPS coordinates from one survey flight and planned a new waypoint flight using our DJI Mavic Pro drone and GS Pro app. The Mavic was ideal for this task as it could lower to ~1 m altitude with much lower assumed risk than the more expensive M300. When the autopilot flew the drone to a waypoint (at ~20 m altitude), we paused the mission and manually lowered the drone directly above the candidate, using the original survey image and prediction box as a guide. If the entire team could confidently eliminate that candidate we would remove it; otherwise, we took two to three pictures for later inspection and proceeded via autopilot to the next waypoint. Since live-feed transmission from the drone provided a lower resolution than its camera, we reviewed the images on our field computer in camp. Any candidates that survived this scrutiny proceeded to the fourth stage: in-person inspection.

3. Results and Discussion

Using the above methodology, we recovered the meteorite <50 m from the best-fit fall line, appearing in the 88th image from the third flight on the first day, with an apparent diameter of 27 pixels and a confidence score of 1: a perfect match. Although we have not yet classified the meteorite, its fusion crust (Figure 4) resembles that of other chondrites. It consists of one 70 g piece, approximately 5 × 4 × 3 cm, with a preferentially smoothed side. Initial CT scans show chondrules along with a mixture of metal and sulfide grains indicative of an equilibrated chondrite. Analysis of our two thick section samples via Scanning Electron Microscope and Electron




Figure 2. Favorable and unfavorable prediction distributions from two images. Given a 70% confidence threshold, Image/Distribution A will return 3 meteorite candidates, while Distribution B will return >100 candidates. B-like images are later used for retraining. For clarity, the number of detections displayed in image B is capped at 50.



Figure 3. The four stage process for eliminating false positives and verifying meteorite candidates. From left to right: (1) Grid GUI. (2) Zoom-pan GUI. (3) Drone visit. (4) In-person visit.



Figure 4. The recovered meteorite as seen in person (top two panels), and from the survey drone (bottom panel). For scale, a 15 cm long felt pen is placed next to the meteorite (top right). The yellow box in the bottom image is 22 cm on one side.

Microprobe, which can definitively determine meteorite type, is forthcoming.

While this meteorite is a great discovery that will hopefully spur the efficient collection of further meteorites, we also made some important discoveries about our methodology. Our data processing rate for both the machine-learning algorithm and manual candidate sorting must be improved in the future, or the length of the trips must be extended as the meteorite appeared on the third of 43 survey flights. We were able to process four flights, totaling 5096 images, which produced 46,501,000 tiles for our algorithm, identifying 56,384 first-stage candidates (Figure 3). Using our candidate elimination process, we produced 259 second-stage candidates, visiting 38 of these as third stages with the Mavic Pro, finally searching for four candidates in person (fourth stage). These four candidates (including the recovered meteorite) are shown in Figure 5 along with four training tiles for comparison. If we were instead required to process all 57,255 images, we would not have been able to complete it on-site, and would have to return on a later trip to follow up on meteorite candidates. That being said, the

meteorite was located <50 m from our “ideal” fall line, which is surprising considering the side-to-side uncertainty in the fall line. With this in mind, we may in the future prioritize searching the area immediately around the ideal fall line.

Before we embark on our next searching trip to one of our 35 unvisited, logged meteorite falls, we must alleviate some logistical bottlenecks in our system. This includes locally networking laptops to our main processing computer on-site to allow more than one team member to sort meteorite candidates from false positives at a time. For each new fall that we visit we will also bring meteorites to retrain our model on generated, local training data to enhance meteorite detection chance and mitigate false positives. In lieu of real meteorites, images of black-painted rocks collected on-site can also work (Anderson et al. 2020).

Although we recovered the meteorite we did not really train a meteorite detection algorithm; instead, we created an anomaly detector, in this case trained for the Nullarbor. During the course of devising this strategy, we encountered false positives such as tin cans, bottles, snakes, kangaroos, and piles of bones




Figure 5. Top: four meteorite tiles selected from our training data. Middle: the recovered meteorite. Bottom: the other three meteorite candidates visited for a second time with the Mavic Pro, and later inspected in person.

from multiple animals. We also notice that when we predict on survey images taken directly over our campsite, the algorithm ferociously identifies our items and equipment, none of which is represented in the training data. We hope that our findings and methodology prove useful for training neural networks in other low-occurrence or anomaly detection problems, such as wildlife monitoring, or search and rescue.

This work was funded by the Australian Research Council as part of the Australian Discovery Project scheme (DP170102529, DP200102073), and receives institutional support from Curtin University. Fireball data reduction is supported by resources provided by the Pawsey Supercomputing Centre with funding from the Australian Government and the Government of Western Australia. The DFN data reduction pipeline makes intensive use of Astropy, a community-developed core Python package for Astronomy (Price-Whelan et al. 2018). We would like to thank Geoff Deacon and Peter Downes at the Western Australian Museum for loaning us meteorites from the collection.

ORCID iDs

Seamus L. Anderson <https://orcid.org/0000-0002-8914-3264>
 Martin C. Towner <https://orcid.org/0000-0002-8240-4150>
 Hadrien A. R. Devillepoix <https://orcid.org/0000-0001-9226-1870>

Eleanor K. Sansom <https://orcid.org/0000-0003-2702-673X>
 Patrick M. Shober <https://orcid.org/0000-0003-4766-2098>

References

- AlOwais, A., Naseem, S., Dawdi, T., et al. 2019, in 2nd Int. Conf. on Signal Processing and Information Security (ICSPIS), 1 (Piscataway, NJ: IEEE)
- Anderson, S. L., Bland, P. A., Towner, M. C., & Paxman, J. P. 2019, in 50th Lunar and Planetary Science Conf. (Houston, TX: LPI), 2426
- Anderson, S., Towner, M. C., Bland, P. A., et al. 2020, *M&PS*, 55, 2461
- Bland, P. A., Spurny, P., Bevan, A. W. R., et al. 2012, *AuJES*, 59, 177
- Chollet, F., Zhu, Q. S., Rahman, F., et al. 2015, Keras Release v2.8.0, <https://github.com/fchollet/keras>
- Citron, R. I., Shah, A., Sinha, S., et al. 2017, in 48th Lunar and Planetary Science Conf. 48 (Houston, TX: LPI), 2528
- Citron, R. I., Jenniskens, P., Watkins, C., et al. 2021, *M&PS*, 56, 1073
- Colas, F., Zanda, B., Vaubaillon, J., et al. 2015, in Proc. of the Int. Meteor Conf. 37 ed. J.-L. Rault & P. Roggemans (Mechelen: International Meteor Organization)
- Cuzzi, J. N., Hogan, R. C., & Shariff, K. 2008, *ApJ*, 387, 102
- Devillepoix, H. A. R., Sansom, E. K., Bland, P. A., et al. 2018, *M&PS*, 53, 2212
- Devillepoix, H. A. R., Sansom, E. K., Shober, P. M., et al. 2022, arXiv:2202.06641
- Devillepoix, H. A., Cupák, M., Bland, P. A., et al. 2020, *P&SS*, 191, 105036
- Gritsevich, M. I., Stulov, V. P., & Turchak, L. I. 2012, *CosRe*, 50, 56
- Howie, R. M., Paxman, J., & Bland, P. A. 2017, *ExA*, 43, 237
- Jansen-Sturgeon, T., Sansom, E. K., & Bland, P. A. 2019, *M&PS*, 54, 2149
- Nakamura, T., Noguchi, T., Tanaka, M., et al. 2011, *Sci*, 333, 1113
- Oberst, J., Molau, S., Heinlein, D., et al. 1998, *M&PS*, 33, 49

- Olech, A., Zoladek, P., Wisniewski, M., et al. 2006, in Proc. of the Int. Meteor Conf., ed. L. Bastiaens et al. (Mechelen: International Meteor Organisation)
- Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *AJ*, **156**, 123
- Sansom, E. K., Bland, P., Paxman, J., & Towner, M. 2015, *M&PS*, **50**, 1423
- Sansom, E. K., Gritsevich, M., Devillepoix, H. A. R., et al. 2019, *ApJ*, **885**, 115
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., & Liu, Z. 2021, A Description of the Advanced Research WRF Model Version 4.3, National Center for Atmospheric Research
- Spurný, P., Borovička, J., & Šrbený, L. 2006, in IAU Symp. 236, Near Earth Objects, our Celestial Neighbors: Opportunity and Risk, ed. G. B. Valsecchi & D. Vokrouhlický (Cambridge: Cambridge Univ. Press), 121
- Towner, M. C., Cupak, M., Deshayes, J., et al. 2020, *PASA*, **37**, e008
- Towner, M. C., Jansen-Sturgeon, T., Cupak, M., et al. 2022, *PSJ*, **3**, 44
- Trigo-Rodríguez, J. M., Llorca, J., Castro-Tirado, A. J., et al. 2006, *A&G*, **47**, 6
- Zender, J., Rudawska, R., Koschny, D., et al. 2018, in Proc. of the Int. Meteor Conf., ed. M. Gyssens & J.-L. Rault (Mechelen: International Meteor Organization)