

JGR Space Physics

RESEARCH ARTICLE

10.1029/2023JA032204

Key Points:

- A technique is developed to generate synthetic head echoes as observed via radar instruments
- Synthetic head echoes are used to train convolutional neural networks to detect head echoes at three radar facilities
- The neural networks identify the majority of real head echoes at each facility, demonstrating the technique's usefulness at any facility

Correspondence to:

T. Hedges,
thedges@stanford.edu

Citation:

Hedges, T., Lee, N., & Elschot, S. (2024). Meteor head echo detection at multiple high-power large-aperture radar facilities via a convolutional neural network trained on synthetic radar data. *Journal of Geophysical Research: Space Physics*, 129, e2023JA032204. <https://doi.org/10.1029/2023JA032204>

Received 20 OCT 2023

Accepted 9 MAR 2024

Meteor Head Echo Detection at Multiple High-Power Large-Aperture Radar Facilities via a Convolutional Neural Network Trained on Synthetic Radar Data

T. Hedges¹ , N. Lee¹ , and S. Elschot¹

¹Department of Aeronautics and Astronautics, Stanford University, Stanford, CA, USA

Abstract High-power large-aperture radar instruments are capable of detecting thousands of meteor head echoes within hours of observation, and manually identifying every head echo is prohibitively time-consuming. Previous work has demonstrated that convolutional neural networks (CNNs) accurately detect head echoes, but training a CNN requires thousands of head echo examples manually identified at the same facility and with similar experiment parameters. Since pre-labeled data is often unavailable, a method is developed to simulate head echo observations at any given frequency and pulse code. Real instances of radar clutter, noise, or ionospheric phenomena such as the equatorial electrojet are additively combined with synthetic head echo examples. This enables the CNN to differentiate between head echoes and other phenomena. CNNs are trained using tens of thousands of simulated head echoes at each of three radar facilities, where concurrent meteor observations were performed in October 2019. Each CNN is tested on a subset of actual data containing hundreds of head echoes, and demonstrates greater than 97% classification accuracy at each facility. The CNNs are capable of identifying a comprehensive set of head echoes, with over 70% sensitivity at all three facilities, including when the equatorial electrojet is present. The CNN demonstrates greater sensitivity to head echoes with higher signal strength, but still detects more than half of head echoes with maximum signal strength below 20 dB that would likely be missed during manual detection. These results demonstrate the ability of the synthetic data approach to train a machine learning algorithm to detect head echoes.

Plain Language Summary Meteors, colloquially referred to as “shooting stars,” are observed when small particles, known as meteoroids, encounter Earth's atmosphere at tens of kilometers per second. The meteoroid surface heats up and emits gas particles as it undergoes high-energy collisions with atmospheric molecules. These neutral gas particles undergo further collisions that cause them to ionize and turn into a plasma surrounding the meteoroid. Radar facilities are capable of observing this plasma. Studying meteor observations tells us about the meteoroid and its origins. In this work, meteor observations are performed at three radar facilities. Radar data can be represented as an image, and a computer vision machine learning algorithm identifies meteors in these images, which would be prohibitively time-consuming if done manually. A method is developed to generate simulated radar images and use them to train the algorithm, avoiding manual identification of thousands of real meteors for the purpose of training. This technique is accurate at all three facilities, and correctly identifies at least 70% of the meteors contained in each data set. This proves that the algorithm works for any radar data that contains meteors, and will allow quick training of machine learning algorithms to detect meteors in future experiments.

1. Introduction

High-power large aperture (HPLA) radar instruments have been used for decades to observe meteoroid particles undergoing hypersonic entry into Earth's atmosphere (Evans, 1965). Many meteors observed via such instruments originate from particles that are micrometers in diameter (Close et al., 2000) and micrograms in mass (Janches et al., 2008), smaller than those observed via sky cameras or optical systems (Brown et al., 2017; Campbell-Brown & Close, 2007), but larger than those observed on space-based impact detectors (Baggaley et al., 2007). HPLA instruments are capable of observing head echoes and nonspecular trail echoes, in addition to the specular trail echoes observed by traditional meteor radars (Stober et al., 2021). By understanding the plasma physics that drives meteor radar observations, one can estimate properties of the parent meteoroid such as mass and bulk density via their radar observations (Close et al., 2005, 2012), quantifying potential hazards to spacecraft and astronauts (Close et al., 2010). Furthermore, meteors act as a probe of the neutral atmosphere through which they

travel, enabling estimation of the atmospheric density (A. Li & Close, 2016; Limonta et al., 2020). Meteors are particularly useful above the altitudes where weather balloons can perform measurements, but below the altitude where satellites in low-Earth orbit experience drag, from which atmospheric density can be estimated (A. Li & Close, 2015).

As a meteoroid undergoes entry at altitudes between 80 and 130 km, its material vapourizes due to sputtering and thermal ablation (Guttormsen et al., 2020; Popova et al., 2001). The resulting neutral particles undergo high-energy collisions with atmospheric molecules traveling between 11 and 73 km per second, depending on the incoming trajectory of the meteoroid relative to Earth (Blanchard et al., 2022), which is sufficient to ionize many of the meteoric and surrounding atmospheric particles. A plasma cap is formed in the immediate vicinity of the meteoroid, which moves with the meteoroid and reflects radio waves (Dyrud et al., 2008; Marshall et al., 2017; Sugar et al., 2021). This produces the head echoes frequently observed at HPLA radar facilities that transmit power on the order of megawatts (Janches & Chau, 2005). As of more recently, non-HPLA radar facilities, such as the Southern Argentine Agile Meteor Radar, have enhanced capabilities that enable head echo observations at sub-megawatt power (Janches et al., 2014; Panka et al., 2021).

As plasma undergoes momentum-scattering collisions with the atmospheric neutral particles and slows down relative to the surrounding atmosphere, it forms a trail along the path traveled by the entering meteoroid. The motion of the plasma relative to the stationary neutrals can cause a Farley-Buneman gradient-drift instability to develop into plasma turbulence (Oppenheim & Dimant, 2015; Oppenheim et al., 2000) and form field-aligned irregularities (FAI) capable of scattering radio waves incident perpendicular to the background magnetic field. This is one mechanism that creates meteor radar signatures known as nonspecular trail echoes (Close et al., 2008; Dyrud et al., 2005). In many cases, nonspecular trails are observed at equatorial radars that point directly zenithward, but non-equatorial radars can also observe these signatures if the beam is pointed appropriately. Non-FAI nonspecular trail echoes have also been observed at higher latitudes (Chau et al., 2014; Kozlovsky et al., 2020), and research efforts to explain these observations are ongoing.

This work focuses specifically on meteor head echo signatures, which are the most frequently observed meteor radar signatures in HPLA radar data. At facilities including Jicamarca Radio Observatory and Millstone Hill Observatory, head echoes can appear thousands of times per hour (Erickson et al., 2001; Y. Li et al., 2020). In general, the head echo observation rate, and the detectability of any particular head echo, depends on factors including facility carrier frequency, power, geographic location and beam pointing direction, in addition to the incoming meteor population at a given time of day and year (Sparks et al., 2009). Taking into account these factors, one can use information from the head echo population as a whole to estimate physical quantities such as the atmospheric neutral density (A. Li & Close, 2016). One can also utilize the observations of head echoes to more efficiently identify trail echoes, since trail echoes frequently appear after head echoes.

Many prior meteor experiments rely on manual searches of radar data to identify head echoes, using heuristics such as received signal power and tools such as Doppler-shifted matched filter banks (Hedges et al., 2022). However, these approaches require one to either search through vast sets of data, which is prohibitively time consuming for most experiments, or only consider signals above a certain threshold. The threshold approach does not differentiate head echoes from phenomena such as sporadic E and the equatorial electrojet, which are extremely common at radar facilities especially near the equator. Furthermore, this approach neglects the weakest head echoes below the signal threshold chosen, effectively neglecting the low end of the meteoroid mass range observable by HPLA radar instruments. To address these issues, Y. Li et al. (2020) developed a rigorous method that calculates the probability that a head echo occurs at every radar pulse and range, without relying on a signal threshold. However, the method is computationally expensive and only performs well for very long coded pulses.

Machine learning methods have previously been developed to automate the identification of head echoes. Initially, a method was developed that applies a modified Hough transform to identify possible head echoes, followed by a convolutional neural network (CNN) to reject false positive identifications (Y. Li et al., 2022). This method minimizes false positives to less than 1% overall, but is computationally expensive and complex to implement. This motivated development of another method that relies purely on a CNN to detect head echoes, and achieves high precision and sensitivity at Jicamarca Radio Observatory (Y. Li et al., 2023).

Despite the success of machine learning methods at Jicamarca Radio Observatory, existing methods rely on the abundance of manually labeled head echoes from pre-existing data collected at the same facility and with the same

experiment configuration. Parameters that influence the appearance of head echoes in raw data, such as the receiver sample rate and chosen pulse code, must be identical between experiments where training data is gathered. As new facilities and experiment types are employed to observe meteors, there remains a need to train such algorithms without relying on pre-existing data. We therefore develop an algorithm to generate synthetic head echoes for any meteor radar experiment, and additively combine them with a subset of other ionospheric phenomena, radar clutter, or noise observed at a particular facility. The resulting examples are used to train a simple CNN architecture. Identifying sufficient examples of data that are not head echoes requires a few hours of manual overhead, but this is much less than the days of effort required to label thousands of head echoes for training. Our CNN architecture accepts raw radar data as input without pre-processing, and therefore is extremely computationally efficient.

In this paper, we first describe the three HPLA radar facilities and data sets used to develop this technique. We describe our CNN detection methodology and architecture, followed by our method to generate synthetic head echoes corresponding to each facility and create training data sets for each facility. We then discuss the performance metrics that result from training a CNN via the synthetic head echo method at each facility and testing it on a subset of real data from the corresponding facility. We compare the properties of head echoes detected via the CNN versus those detected via an exhaustive manual search, and conclude with a discussion of how this work will facilitate improved analysis in future radar experiments that detect head echoes.

2. Radar Experiment and Facilities

On October 10th and 11th, 2019, a meteor radar data collect was performed concurrently for 8 hr total at three radar facilities, each with a co-located transmitter and receiver. The facilities include Resolute Bay Incoherent Scatter Radar North (RISR-N), Millstone Hill Observatory (MHO), and Jicamarca Radio Observatory (JRO), with the intent to study neutral density variation in the lower thermosphere via meteor observations at varying latitudes but similar longitudes. Each facility is unique in its operational parameters and geographic location, which leads to observational biases in the meteor population. Biases inherent to latitude include the beam orientation and its angle relative to meteoroid sources, varying geomagnetic field orientation and ionospheric conditions. Other biases can result from the radar carrier frequency, incident power (Urbina & Briczinski, 2011), beam pattern, incident polarization, and local solar time. One must therefore consider these biases in developing a machine learning algorithm that detects head echoes.

Each facility transmits long phase-coded pulses to maximize transmitted signal power and thus increase sensitivity. Long pulses take advantage of the sparsity of meteor radar observations, since there is usually only a single head echo being observed at any given time and range (Volz & Close, 2012). Phase-coded pulses allow Doppler shifted matched filters to be applied to the raw data in the vicinity of the head echo to minimize its range ambiguity. The autocorrelation functions of long pulse codes contain range sidelobes surrounding a central peak, which are exacerbated by any mismatch in Doppler shift of the matched filter. However, assuming the Doppler shift is determined correctly, and since separate head echoes rarely overlap in such a way as to interfere with one other, the central peak clearly indicates the precise range of the meteor despite the presence of range sidelobes.

The RISR-N and JRO facilities transmit a minimum-sidelobe 51-baud (MSB 51) code of duration 51 μ s. MHO transmits a Barker-7 code of length 42 μ s, since the facility was not capable of a baud length as short as 1 μ s at the time of experiment, unlike the other facilities. This code produces nearly comparable signal-to-noise (SNR) performance to that of the MSB 51 code. The RISR-N facility receives signal at a sampling rate of 2 MHz, enabling range resolution of 75 m, whereas JRO and MHO receive at 1 MHz, enabling range resolution of 150 m. The JRO radar beam is directly zenith-pointing, and the RISR-N beam is almost zenith-pointing, allowing for shorter interpulse period to reach the meteor altitude range of 80–120 km. The MHO beam points at a 45° elevation directly west to avoid ground clutter that would otherwise occur at the location. Each facility transmits at interpulse periods of milliseconds to ensure head echoes are resolved across many pulses and minimize uncertainty in the measurement of range and its evolution with time. Experiment parameters and details for each radar facility are provided in Table 1.

The dynamics of meteoroids entering Earth's atmosphere and their resulting head echoes are dependent on the conditions of the lower thermospheric neutral atmosphere and ionosphere during the experiment. Table 2 specifies the solar and geomagnetic indices for the dates and times of experiment. The indices reflect a quiet geomagnetic day, as expected for the 2019 solar minimum.

Table 1
Parameters for Each Radar Facility and Experimental Configuration

Parameter	RISR-N	MHO	JRO	Units
Longitude	−94.91	−72.47	−76.87	deg
Latitude	74.73	44.16	−11.95	deg
Observation date	October 10–11, 17–18	October 10–11	October 10–11	–
Observation time	09:00–13:00	09:00–13:00	05:00–13:00	UTC
Beam Azimuth	26	270	–	deg
Beam Elevation	86	45	90	deg
Carrier Frequency	442.5	442.9	49.9	MHz
Pulse Code	MSB 51	Barker 7	MSB 51	–
Transmit Power (peak)	2	2	6	MW
Baud Interval	1	6	1	μs
Interpulse Period	1.4	2.0	1.25	ms
Sample Rate	2	1	1	MHz
Noise floor (raw data) ^a	13.98	29.45	29.83	dB
Noise floor (matched filter)	35.19	47.48	48.00	dB

^aIn the 4 hr of this experiment, it was observed that the noise floors at each facility vary no more than 1 dB, enabling the use of constant average noise floors.

3. Machine Learning Method

To parse through the many hours of radar data and identify head echoes without relying on days of manual effort, a CNN technique is developed to perform this task automatically. To keep the technique simple and reduce the necessary computational power, the CNN is trained to determine whether small segments of raw data contain a head echo. A simple CNN architecture with three layers is chosen. By simplifying the meteor trajectory dynamics and modeling the reflection that generates the head echo signature, synthetic head echoes are generated to train the CNN. To ensure the CNN appropriately differentiates between head echoes and other phenomena, examples of data with such phenomena or noise are combined with the synthetic head echoes to generate training examples. This allows the CNN to be trained using tens of thousands of training examples. The CNN technique is implemented in Python using the PyTorch library (Paszke et al., 2019). Since PyTorch supports CUDA functionality, the CNN training and testing calculations are performed on a GPU for accelerated performance.

3.1. Head Echo Detection Approach

We present a CNN architecture that accepts raw radar data as input, with no pre-processing beyond separating the data into segments of consistent size. This enables extremely computationally inexpensive classification of each segment. A matched filter is ill-suited for pre-processing since it requires prior knowledge of the Doppler shift for each head echo. One can use a technique based on the fast Fourier transform (FFT) that algorithmically accounts

Table 2
Solar and Geomagnetic Indices on the Experiment Dates and Times

Index	10 October 2019			11 October 2019		
Time (UTC)	09–12	12–15	Full-day average	09–12	12–15	Full-day average
K_p	2.00	1.33	2.42	2.00	1.33	1.71
A_p	7	5	13.4	7	5	6.9
$F_{10.7}$	–	–	67.5	–	–	68.5

Note. The low $F_{10.7}$ values reflect the solar minimum at the time of experiment, and the low values of K_p and A_p reflect a quiet geomagnetic day (Matzka et al., 2021).




Figure 1. Raw data is split into segments with some overlap, and the convolutional neural network classifies segments as positive (label of 1) or negative (label of 0) depending on whether a head echo is present. In this diagram (not to scale), red lines indicate segment boundaries, yellow slanted lines indicate head echoes, and yellow clouds indicate nonspecular trail echoes that follow some head echoes. The segment is considered positive if any part of any head echo is detected. For this work, trail echoes are not classified as positive.

for the Doppler shift to decode the raw data, known as FFT decoding, but this method is computationally expensive (Y. Li et al., 2023).

HPLA radar raw data effectively consists of a two-dimensional array of complex in-phase and quadrature (I&Q) voltage values, in some cases with additional channels that capture dual polarization or interferometric measurements from the received signal (Chau & Woodman, 2004; Close et al., 2011). The CNN architecture accepts the real and imaginary components of the raw data as separate image channels of the input. When the data is visualized as a range-time intensity (RTI) image, as shown in the radar image plots throughout this paper, the horizontal axis of the data represents radar pulses (i.e., time), and the vertical axis represents range gates (i.e., altitude). The decibel power, P_{dB} , is computed via

$$P_{dB} = 20 \log_{10} |c_{rx}| - n_f, \quad (1)$$

where c_{rx} is the raw data or decoded data via a matched filter at a particular pulse and range gate, and n_f is the decibel noise floor as specified in Table 1 for this experiment.

The raw radar data arrays contain hundreds of range gates and thousands of consecutive pulses at each facility, so the arrays are broken into separate segments to be analyzed by the CNN. We use segments with size of 150 pulses by 150 range gates for MHO and JRO. Since RISR-N uses a sample rate twice as large as the other two facilities, we use segments of 150 pulses by 300 range gates for this facility. The segment sizes are chosen to ensure the majority of a particular head echo is generally captured within the segment, given the pulse duration and sample rates used at each facility, without being so large as to consume excessive GPU memory usage when training the CNN. Each segment overlaps the adjacent segments on all four sides by 30 pulses or range gates. Since some head echoes can occur for only a few pulses, the overlap on the left and right segment sides ensures that such short-duration head echoes are fully captured within at least one segment, increasing the ability of the CNN to detect said head echoes. The overlap in range gates on the bottom and top segment sides is not necessary, given the use of long coded pulses that are spread across many range gates, but is included for simplicity. If more than one head echo is contained within a segment, the segment should be classified as positive. For this work, if a trail echo or other meteor signature is contained in a segment, the segment should be classified as negative. Figure 1 demonstrates how radar data is split into chunks and then labeled where head echoes are present, assuming perfect image classification. The data contained within a bounding box that surrounds adjacent positive classifications can then be fed into a head echo analysis routine to determine the range, velocity, and other such properties of every head echo within the region.

3.2. Convolutional Neural Network Architecture

Our CNN model architecture is inspired by the ResNet architecture (He et al., 2015) and includes three convolutional layers followed by a fully connected layer, as depicted in Figure 2. This architecture was chosen since it outperformed an architecture with two layers, but adding a fourth layer did not continue to improve performance. The first two convolutional layers employ a 3-by-3 convolution operation, followed by the leaky rectified linear




Figure 2. Top: Diagram of the convolutional neural network architecture. Bottom: Diagram of an individual convolutional layer. The first and second layers in our architecture are exactly as depicted, whereas the third layer replaces the 2-by-2 maxpool with a global maxpool.

unit (ReLU) activation function and 2-by-2 maxpool operation. The leaky ReLU activation function is taken elementwise over the output of the convolution operation and is defined as

$$f(x) = \begin{cases} ax, & x < 0, \\ x, & x \geq 0, \end{cases} \quad (2)$$

where $a = 0.1$ in this CNN architecture. The 2-by-2 maxpool operation splits each image channel into 2-by-2 blocks and then keeps only the largest element, which scales down the image channels, as visualized in Figure 2. Although the maxpool operations scale down the image at each layer, the number of image channels increases throughout the network, since each convolution operation contains multiple kernels that operate in parallel and produce output on a separate layer (15 kernels on layer 1, 50 on layer 2, and so on).

The third and final convolutional layer replaces the 2-by-2 maxpool with a global maxpool, which keeps only the maximum element from each image channel and reduces the layer output to a one-dimensional vector with these elements. This is followed by a fully connected layer and softmax operation, which reduces the vector to a smaller output vector of length 2, where each component is the probability that the input image has a label of 0 or 1. In general, the softmax operation yields a vector with K components, where K is the number of labels that an image can be classified by. The i th component is then determined by the equation

$$\sigma(\vec{y})_i = \frac{e^{y_i}}{\sum_{j=1}^K e^{y_j}}, \quad (3)$$

where $\vec{y} \in \mathbb{R}^K$ is the output from the fully connected layer. In this case, since $K = 2$, the value of σ_2 specifies the probability that the raw data segment contains a head echo, as determined by the CNN. If this probability is greater than 0.5, the segment is classified as positive. The CNN architecture contains a total of 142,987 coefficients, or weights, that are optimized during training.

3.3. Synthetic Head Echo Model

To generate synthetic head echoes, we assume the head echo range is exponential in time, since it travels through the atmosphere with exponential density dependence on altitude, and the drag force experienced by the meteoroid is proportional to atmospheric density. The range is therefore given by

$$r(t) = \left(r_0 - \frac{v_0^2}{a_0} \right) + \frac{v_0^2}{a_0} e^{-(a_0/v_0)t}, \quad (4)$$

where r_0 is the initial range and v_0 is the initial radial velocity, which is the velocity component along the radar beam, where the value is defined positive when the velocity vector points toward the facility. The value of a_0 is the initial radial deceleration, which is the acceleration projected along the beam direction, defined positive when this vector points away from the facility (i.e., a deceleration as the meteor descends). The initial values are taken at the time of first head echo observation.

Although these assumptions about the meteoroid dynamics neglect more complex physical phenomena such as differential ablation (Dyrud & Janches, 2008; Janches et al., 2009) and fragmentation, and the dual polarization scattering that may result from fragmentation events is not included in the synthetic model (Close et al., 2011), this method still generates head echoes that closely resemble head echoes observed in real data. In practice, large portions of real head echoes will resemble those in the synthetic training set, allowing the CNN to detect them despite their more complex behavior.

To generate the raw radar data corresponding to a meteoroid, we assume the received pulse is created via a single-body Doppler-shifted reflection of the transmitted pulse. For convenience, we define the variable $\tilde{t} = t - t_i$, where $t_i = i\Delta t_p$ is the time at which the i th pulse starts transmitting, and Δt_p is the interpulse period (i.e., pulse resolution). It is assumed that the magnitude of the received signal varies solely due to transverse meteor motion through the radar beam, as opposed to variation of the meteor plasma density or shape throughout the time of observation. Therefore, we assume the signal strength over time is proportional to the main lobe of the sinc function, $\text{sinc}(s) = \frac{\sin(\pi s)}{\pi s}$, where $s = \frac{2\tilde{t}}{t_f} - 1$ and t_f is the time of the final pulse that the head echo is observed. More complex formulations of SNR were tested, such as including multiple local SNR maxima via a Fourier sine series throughout the duration of observation, but this approach did not produce superior results to a simple sinc function for each head echo. The peak SNR in decibels is given by the parameter $S_{\max}$. Note that this SNR value is defined as the head echo signal strength in raw data, which differs from the SNR definition used to quantify signal strength of real head echoes, which is the SNR in decoded data after a matched filter with the appropriate Doppler shift is applied.

The function $c_{tx}(\tilde{t})$ represents the complex voltage signal of the transmitted pulse and is nonzero when $\tilde{t} \in [0, B(L-1)]$, where L is the pulse length in bauds and B is the baud interval. The complex raw data of i th received pulse of the head echo is then determined by the equations

$$c_{rx,i}(\tilde{t}) = A \text{sinc}\left(\frac{2t_i}{t_{L-1}} - 1\right) c_{tx}\left(\tilde{t} - \frac{2r(t_i^*)}{c}\right) e^{i\phi(t)}, \quad (5)$$

$$\phi(t) = \frac{4\pi f}{c} \int_0^t v(t') dt' + \phi_0, \quad (6)$$

$$A = 10^{\frac{1}{20}(S_{\max} + n_f)}, \quad (7)$$

where f is the facility carrier frequency, c is the speed of light, n_f is the noise floor of raw data in decibels, and t_i^* is the time when the i th transmitted pulse reflects from the head echo. The values of f and n_f are defined in Table 1 for each facility. The phase parameter ϕ_0 represents the phase of the received complex signal at the time of first observation. The value of t_i^* can be determined exactly by solving the implicit expression

$$c(t_i^* - t_i) = r(t_i^*), \quad (8)$$

Table 3
Parameters Determined Randomly for Each Head Echo During Synthetic Data Generation

Parameter	Meaning	Minimum value	Maximum value	Units
r_0	Initial range	75	125	km
v_0	Initial radial vel.	11	73	km/s
a_0	Initial radial decel.	0.01	55	km/s ²
ϕ_0	Initial signal phase	0	2π	rad
t_{L-1}	Observation time	$12\Delta t_p$	$60\Delta t_p$	s
$S_{\max, \text{RISR-N}}$	Peak SNR (RISR-N)	0	8	dB
$S_{\max, \text{MHO}}$	Peak SNR (MHO)	1	14	dB
$S_{\max, \text{JRO}}$	Peak SNR (JRO)	0	20	dB

Note. Δt_p is the interpulse period for the corresponding radar facility for which synthetic data is generated, and the signal-to-noise (SNR) values define the head echo strength in raw data, which is smaller than the SNR after a matched filter is applied.

but it is sufficient to approximate t_i^* to simplify the implementation of synthetic data generation. The simplest such approximation is to assume the meteor does not move much within one pulse, in which case

$$t_i^* \approx t_i + \frac{r(t_i)}{c}. \quad (9)$$

This approximation is improved by instead assuming that the meteor does not decelerate much during the time that the radio waves travel,

$$t_i^* \approx t_i + \frac{r(t_i) + v(t_i)(t_i^* - t_i)}{c}, \quad (10)$$

and then plugging in $t_i^* - t_i$ via the simpler approximation of Equation 9. Hence, the following equation is used to approximate t^* for the purpose of synthetic head echo generation:

$$t^* \approx t_i + \frac{r(t_i) + v(t_i)(r(t_i)/c)}{c}. \quad (11)$$

All of the parameters in Equations 5–7 are determined randomly for each synthetic head echo via a uniform probability distribution within the ranges specified in Table 3.

Although a more traditional CNN architecture is trained for this work, the synthetic head echo model can be adapted to train any neural network model, including more advanced architectures such as the You Only Look Once (YOLO) algorithm, which is capable of identifying objects and their bounding boxes (Bochkovskiy et al., 2020).

3.4. Generation of Training Examples

A large set of head echoes can easily be generated via the synthetic model. However, one must position these head echoes in an image segment of the appropriate size and combine them with background noise or phenomena in order to generate training data examples. A set of 10,000 synthetic head echoes at JRO, and 25,000 head echoes at RISR-N and MHO, are generated with randomized parameters for each facility via the aforementioned procedure. The bounding box of each head echo is positioned at a random floating-point coordinate within an otherwise empty raw data segment, with the constraint that no more than 25% of any side length of the bounding box around the head echo extends from a segment edge. The x -coordinate of the first pulse of the head echo is rounded to an integer to align with the pulse resolution. The y -coordinate of the first sample is not rounded to an integer; instead the simulated head echo samples are linearly interpolated to the exact coordinates of the range gates of the

segment. Physically, this accounts for the fact that the received data need not be sampled in alignment with the individual bauds of the signal; there will be some arbitrary delay.

Once the head echo has been placed in its segment, background phenomena must be added to make the raw data realistic and ensure the neural network learns to differentiate between head echoes and other phenomena. Unfortunately, it is currently not easy to generate background phenomena synthetically, since there is a wide range of unique phenomena observed by any given facility. Although models have been constructed for nonspecular trails and the equatorial electrojet (Dyrud et al., 2005; Hysell et al., 2002), attempting to simulate radar data for each phenomena would be prohibitive. Thus, we combine synthetic head echoes with examples of background phenomena obtained from the real raw data that do not contain head echoes, which will henceforth be referred to as real negative segments. Real negative segments are much more common and easier to identify than segments that contain head echoes.

We initially identify more than 2,000 real negative segments of raw data, which requires a few hours of manual overhead for each facility. These segments are carefully selected to ensure they are truly negative; that is, contain no part of any object that resembles a head echo. The negative set chosen from JRO includes many instances of nonspecular trail echoes and the equatorial electrojet (EEJ). The MHO set contains some clutter and potential sporadic E events, although not as many events as JRO. Unlike the other facilities, RISR-N rarely observes objects other than head echoes throughout the experiment duration, so the negative segments chosen for RISR-N are simple noise examples. To reduce manual overhead as much as possible, we employ augmentation techniques to artificially increase the number of real negative segments at each facility from approximately 2,000 to 20,000. Such techniques include taking arbitrary linear combinations of the selected negative segments, and arbitrarily translating the segments in the horizontal and vertical directions, wrapping data around the edges of the segment. Although this process results in nonphysical discontinuities, it does not bias the training process in any way, since such discontinuities are not indicative of whether head echoes exist. A negative segment is then randomly selected to be additively combined with each synthetic head echo to create the set of up to 25,000 positive training examples. For each positive training example, the corresponding negative segment without the head echo is retained as part of the training set, for a total of up to 50,000 training examples. The process of combining a synthetic head echo with a real negative segment is demonstrated in Figure 3, along with examples of training data generated for JRO.

3.5. Neural Network Training Process

A single neural network per facility is trained on segment examples containing synthetic head echoes generated for the corresponding facility. The same CNN architecture is used at each facility, but the resulting training weights vary based on the experiment parameters and training set.

An algorithm based on mini-batch stochastic gradient descent, known as the Adam algorithm (Kingma & Ba, 2017), is used with the cross-entropy loss function to train the network. Each optimization step takes into account 100 training examples. This batch size was chosen since it is large enough to ensure quick convergence, but small enough to not consume excessive memory during training. The learning rate, which controls step size of the optimization algorithm, starts at 0.001, and then is decreased manually between epochs to a final value of 1×10^{-5} . Since the training set is large, the algorithm converges within only a few epochs, or loops through each training example. Weight decay, which is similar to L2 normalization of the training weights, is employed to protect against overfitting, with a coefficient of one-tenth of the learning rate.

4. Neural Network Performance and Discussion

The CNN is tested on a set of manually labeled raw data segments taken at each facility from the concurrent radar experiment. At RISR-N and MHO, one test set is considered per facility. At JRO, we consider performance against two test sets; one before dawn when the EEJ is not present, and one during dawn where the EEJ is present. In real raw data, we manually search for and identify head echoes using either FFT decoding or a bank of Doppler-shifted matched filters, since head echoes are not always readily visible in raw data images, but it is desired that the CNN detect even the weakest head echoes in raw data. When using these methods, objects are occasionally present that resemble streaks but are not clearly head echoes. Segments containing such objects are considered inconclusive and excluded from calculations of performance metrics.




Figure 3. Realistic segments used to train the convolutional neural network (CNN) are generated by additively combining synthetic head echoes with real segments of ionospheric phenomena or noise from a given facility. The head echo appears range-aliased in raw data due to the long code length, but becomes a streak after a Doppler-shifted matched filter is applied. Note that the CNN is trained on the raw data. A selection of training segments generated for Jicamarca Radio Observatory, along with their labels, are shown. For each example generated with a synthetic head echo, a second example is generated without the head echo but with the same background data.

For each test set, labels for each segment are predicted via the trained neural network, and then checked against the truth labels. Performance metrics are quantified for each test set, including classification accuracy, which is the overall fraction of correct classifications, the precision p , which is the fraction of positive classifications that are correct, the sensitivity s , which is the fraction of positively labeled segments that are correctly identified, and the F_1 score, which is defined as

$$F_1 = 2 \left(\frac{ps}{p + s} \right). \quad (12)$$

The sensitivity metric is broken into two separate categories; sensitivity to head echoes with a post-matched filter maximum SNR of greater than or equal to 20 dB, and sensitivity below this threshold. Head echoes of post-matched filter signal strength below 10 dB are rarely observed in this experiment configuration, so the latter category represents sensitivity to the weakest detectable head echoes. For each test set, the number of test segments with each label, and performance metrics, are provided in Table 4.

RISR-N has the greatest classification accuracy, precision and F_1 score, as expected given the infrequency of objects besides head echoes at this facility. Its reduced sensitivity as compared to MHO, although still high overall, could be a result of the larger segment size used or the more complex pulse code. Since MHO has some sporadic E and cluttering phenomena, its precision is slightly lower than that of RISR-N. JRO has the lowest precision, overall accuracy and F_1 -score, which results from its frequent and prominent observations of objects other than head echoes, especially nonspecular trail echoes which fill larger regions of data and frequently overlap head echoes. The precision drops further when the EEJ is present, but it is still greater than 0.5. This precision is still reasonable, since only one false positive needs to be rejected per true head echo.

Table 4

Number of Segments in Each Test Set, Including Breakdown of Frequency of Each Truth Label (Positive if a Head Echo Is Present, Negative if Not), and Performance Metrics for the Convolutional Neural Network at Each Radar Facility, Including Resolute Bay Incoherent Scatter Radar North (RISR-N), Millstone Hill Observatory (MHO), and Jicamarca Radio Observatory (JRO)

	RISR-N	MHO	JRO (no EEJ)	JRO (EEJ)
Test set size (total)	11,868	7,140	4,836	4,836
Positive segments	157	148	290	122
Negative segments	11,631	6,848	4,373	4,574
Inconclusive segments	80	144	173	140
Classification accuracy	0.996	0.994	0.962	0.974
Precision	0.967	0.881	0.664	0.503
Sensitivity (overall)	0.758	0.808	0.774	0.713
Sensitivity (≥ 20 dB)	0.977	0.970	0.926	0.781
Sensitivity (< 20 dB)	0.538	0.649	0.562	0.543
F_1 score	0.850	0.843	0.715	0.590
Precision via FFT decoding and 20 dB threshold	1.000	0.293	0.259	0.073
Sensitivity via FFT decoding and 20 dB threshold	0.459	0.411	0.581	0.902

Note. At JRO, two separate test sets are evaluated; one without the equatorial electrojet (EEJ), and one with the EEJ present after dawn. The overall sensitivity and precision is also calculated for the method of assuming any segment with signal strength greater than 20 dB contains a head echo, after applying FFT decoding to each raw data segment. This demonstrates how using the CNN for identification improves upon using a signal strength threshold, which is a more traditional method of searching for head echoes.

The CNNs demonstrate greater than 0.7 overall sensitivity to head echoes at each facility, including at JRO when the EEJ is present. Sensitivity is significantly higher for stronger head echoes, as one would expect, whereas for the weaker head echoes, it remains above 0.5 across all facilities. Therefore, one can employ this algorithm to obtain a comprehensive head echo data set, including numerous weaker head echoes, at any facility. Sensitivity to strong echoes decreases when the EEJ is present at JRO since the echoes are often obscured by the EEJ, but most of the head echoes identified via manual searching are still detected. The strong performance at each facility demonstrates that one could apply the synthetic head echo technique at any facility for a future head echo experiment and quickly identify head echoes without an intensive training process.

Prior work for head echo detection at JRO has demonstrated precision and sensitivity values above 0.9 (Y. Li et al., 2023). Although our CNN does not achieve this performance at JRO, it is likely that the reduced precision results from the use of purely synthetic data to train the CNN, in addition to the simplicity of the CNN architecture and lack of pre-processing. It is important to note that the reduced overall complexity enables significantly faster computational performance and ease of implementation. Future work will investigate incorporation of varying proportions of real data into the training set, in addition to more advanced architectures and inclusion of pre-processing to increase precision and sensitivity.

To quantify how the CNN algorithm improves over traditional (i.e., non-machine learning techniques) used to detect head echoes, the performance of the method is compared with that of using a signal threshold to search for head echoes. The signal threshold was performed as follows: FFT decoding is used to decode each segment of raw data, and then segments with any signal above a 20 dB threshold are labeled positive. Such a method is frequently used as a first pass to search for head echoes, but is expected to perform poorly when high-SNR phenomena other than head echoes are prevalent, and is furthermore expected to ignore any head echoes below the threshold signal strength. The overall sensitivity and precision of the threshold method are included in Table 4. Although perfect precision is achieved at RISR-N due to the infrequency of non-head echo phenomena at this facility, the precision drops significantly at MHO and JRO, and drops further when the EEJ is present. Since many head echoes are weaker than the signal threshold, the sensitivity of the method is poor in each test set besides that of JRO with the EEJ present. This occurs because with the presence of the EEJ, many segments are positively labeled despite the




Figure 4. Class activation maps for some head echoes observed amidst other phenomena at the Jicamarca Radio Observatory (JRO) and Millstone Hill Observatory (MHO) facilities. The contour lines indicate which regions of the image the convolutional neural network deems most relevant for positive classification.

technique being insensitive to the actual appearance of a head echo. Use of the CNN for head echo identification is therefore far superior to a signal threshold search, which effectively reduces to a manual search when the EEJ or other prevalent phenomena are present.

To demonstrate the ability of the CNN to identify head echoes and discriminate them from other observations, a class activation map (CAM) can be generated, which is an image that indicates which regions of a CNN input trigger a positive classification. The CAM is generated via intermediate output from any CNN architecture that utilizes a global pooling operation after the convolutional layers. We forward propagate a particular input example into the CNN, and consider the output from the activation function of the third layer before the global maxpool is performed, of dimension 300-by-38-by-38. This output contains 300 channels of 38-by-38 images that represent specific combinations of features in the image that the CNN has detected. Each of the 38-by-38 channels is multiplied by its corresponding scalar weight in the final fully connected layer of size 300, and the resulting arrays are summed together. The values in this array are interpolated to the coordinates of the larger input image, effectively upscaling it to the input resolution, and forming a two-dimensional image that specifies which regions of the input example most influence the CNN to consider it a positive detection. A contour map is generated for this image and then overlaid onto an RTI image generated from the raw data segment. Some CAMs for head echoes correctly identified near other phenomena at JRO and MHO are shown in Figure 4.

From the CAMs, it is clear that the CNN is capable of discerning the head echo signature, even in cases where background phenomena or clutter is present. At JRO, head echoes that are embedded within the equatorial electrojet or precede a nonspecular trail are clearly detected independently. At MHO, brief range-spread clutter is frequently detected, and the CNN detects head echoes separately from such events, even in cases where the clutter is stronger. RISR-N detects cluttering events so infrequently that there is no practical need to eliminate false positives that occur due to such events.

Within the test set corresponding to each facility, with the exception of the test set at JRO where the EEJ is present, each of the head echoes is analyzed to determine its altitude, radial velocity, and SNR. FFT decoding is used to directly solve for an initial approximation of the average Doppler velocity of each head echo, as described in Y. Li et al. (2023). A matched filter with the corresponding Doppler shift is then used to decode the raw data. Within the region of the head echo, the raw data value at the range gate of maximum SNR is chosen for each pulse, and the pulse-to-pulse Doppler shift is determined as described in Hedges et al. (2022) to obtain the most accurate radial velocity throughout the duration of detection. The head echo properties for each facility are presented in the scatter plots in Figure 5. Head echoes that are contained in at least one correctly labeled raw data segment are marked as identified by the CNN, whereas head echoes not correctly identified in any raw data segments are marked as not identified. Note that the fraction of head echoes correctly detected by the CNN is greater than or equal to the overall CNN sensitivity at each facility, since in this case, the sensitivity is defined as the fraction of segments containing head echoes that are correctly identified, and a head echo need only be detected in one raw data segment to be considered detected in the scatter plots.




Figure 5. Scatter plots of head echoes detected at the three facilities, where the vertical and horizontal axes represent altitude and radial velocity, respectively, and the colorbar represents maximum SNR throughout the detection timespan after a matched filter is used to decode the raw data containing the head echo. Circles indicate head echoes where at least one segment containing the head echo was correctly identified by the convolutional neural network, whereas triangles indicate head echoes where no segment containing it is detected by the neural network.

In each of the scatter plots, it is further demonstrated that head echoes with higher maximum signal strength are preferentially detected, but that many weaker head echoes below the 20 dB threshold are also detected. These are head echoes that likely would not be included in an analysis via more traditional detection techniques such as SNR thresholding, at least without significant manual overhead. Therefore, the CNN technique yields a more comprehensive population that includes head echoes resulting from meteoroids that travel farther from the center of the beam, in addition to those of smaller initial mass. Where interferometric measurements are available, such lower-mass observations can be directly identified.

Fewer head echoes are present in the test sets at RISR-N and MHO than at JRO, despite the larger number of raw data segments in the RISR-N and MHO test sets, since these facilities observe head echoes more infrequently than at JRO (Hedges et al., 2022). There are no clear trends amongst altitude or radial velocity in the undetected head echoes at RISR-N or JRO. At MHO, the CNN tends to overlook head echoes more frequently at lower radial velocities. This occurs because the MHO beam was pointed at a 45° elevation angle due west for this experiment, unlike at the other facilities with near-vertical beams. This enables observation of head echoes with near-zero or negative radial velocities; for example, a meteoroid entering due west at 45° from the horizontal will traverse the radar beam perpendicular to its pointing direction, resulting in zero radial velocity. However, the CNN was trained only on head echoes with radial velocity between 11 and 73 km per second. An attempt was made to train the CNN with synthetic head echoes between −20 and 73 km per second for MHO, but the inclusion of synthetic head echoes with near-zero radial velocity caused a significant increase in the false positive rate. This is likely due to similarity between head echoes with little radial motion and therefore little Doppler shift, and non-head echo phenomena including the range-spread clutter frequently observed at MHO, as observed by the CNN.

The radar observations at all three facilities were not coordinated to observe any known meteor showers, so it is assumed that the vast majority of observed meteors are sporadic (Schult et al., 2018). However, interferometric capability is necessary to conclusively confirm or deny the presence of any meteor showers.

5. Conclusions and Future Work

In this paper, a technique to generate synthetic head echo observations given any radar facility and experiment parameters is developed. Furthermore, it is proven that a CNN can successfully identify head echoes in raw radar data without any pre-processing, which greatly reduces computational cost. The utility of the synthetic head echo technique is demonstrated via its ability to train CNNs to detect head echoes at multiple HPLA radar facilities, including RISR-N, MHO, and JRO, which performed concurrent observations in October 2019. Each CNN is tested against a selection of real data from each facility, where raw data segments are manually labeled as either positive (contains one or more head echoes) or negative (does not contain a head echo). Each CNN demonstrates overall sensitivity greater than 0.7, indicating that most head echos present in the data are identified.

The synthetic data approach is valuable because it significantly reduces the manual overhead necessary to train a computer vision algorithm to identify head echoes. Manually labeling head echoes would require days of effort and extensive head echo data from similar experiments (i.e., similar pulse code and other experiment parameters), which is not always available. With the synthetic data approach, one must still manually identify a set of raw data segments that do not contain head echoes, ideally with a comprehensive set of ionospheric phenomena and clutter observed by a facility. Since such segments are much more frequent than those with head echoes, and data augmentation techniques are used to allow fewer such segments to be identified, this process takes hours rather than days.

This technique will enable a more in-depth analysis of head echoes throughout the full data sets at all three facilities that performed concurrent observations. Previous analyses from Hedges et al. (2022) only included hundreds of head echoes at each facility. A more complete analysis containing many more head echoes will elucidate any temporal variation in the head echo population throughout the duration of observations, allowing quantification of variation in the neutral density within each latitude region.

Future work will enhance the capability of the technique to generate physically accurate synthetic data. Incoming meteoroids will be simulated along three-dimensional trajectories that pass through the known antenna pattern of a facility, which produces range and signal strength modulations in head echoes that are not currently captured by the technique. Simulation of the trajectory will further enable a model of the additional receiving raw data channels available at JRO, including dual polarization in the four separate receiving quadrants of the facility. This allows interferometric measurements to be utilized to deduce the location of scattering targets observed by the radar, and thus could allow a neural network to differentiate between head echoes and phenomena such as the EEJ that may not be co-located within the beam. This will enhance the sensitivity of the CNN to head echoes observed at the same range as the EEJ. This approach can be extended to other facilities capable of interferometric measurements to enhance detection, including the Middle Atmosphere Alomar Radar System in Norway (Schult et al., 2013), the Southern Argentine Agile Meteor Radar (Janches et al., 2014), and the middle and upper atmosphere (MU) radar in Japan (Kastinen & Kero, 2022). Additionally, future work will explore integration of real data to the training sets in conjunction with synthetic data, since this may improve performance while still leveraging the benefits of synthetic data.

Subsequent research will explore adaptation of the synthetic data technique to more advanced CNN architectures. In particular, incorporating pre-processing via decoding of raw data could increase the precision and sensitivity to be comparable to what was achieved in previous work. Addition of the capability to identify a bounding box surrounding each head echo in addition to classifying whether a head echo is present will enable complete automation of head echo analysis, including determination of the meteoroid velocity and deceleration versus time (Y. Li et al., 2023). This will enable usage of the technique to completely automate processing of head echoes, including determination of meteoroid velocity and deceleration versus time. Future research will also quantify sensitivity biases of CNN methods including this one, and determine how the identified head echo populations differ from those obtained manually. One must consider that manual identification may still yield a biased set of head echoes, since the weakest and shortest-duration head echoes are easily missed by the human eye. To avoid such bias, one must carefully identify an exhaustive head echo set that includes the weakest and shortest head echoes, as was done for the test sets in this work. Therefore, continued research of machine learning identification techniques will prove critical in the ongoing study of meteor head echoes and their scientific value.

Data Availability Statement

The Python implementation of synthetic data generation and the convolutional neural network algorithm, along with all of the data used to generate the results from this paper, is available on the Stanford Digital Repository at <https://doi.org/10.25740/zn070dk3564> (Hedges et al., 2023). The archive includes each test set from each facility containing the raw data segments used to produce the results of this paper, the tabulated post-matched filter SNR of each head echo in each test set, properties of each head echo versus time used to generate the scatter plots, the examples of clutter or noise needed to generate synthetic data, and the training weights for the CNN used to generate the results of this paper.

Acknowledgments

The authors graciously thank those that provided assistance with data collection: Marco Milla and Karim Kuyeng Ruiz with Jicamarca, Phil Erickson and Ryan Volz with Millstone Hill, and Roger Varney with RResolve Bay. The authors also thank Meers Oppenheim and Glenn Sugar for their advice with data analysis and signal processing. This work, including radar operations and analysis at AMISR RResolve Bay and Jicamarca, was funded by NSF Grant AGS-1920383, and student funding was supplemented by NSF Grant AGS-2048349. Radar operations and analysis at Millstone Hill were supported by NSF grant AGS-1952737 to the Massachusetts Institute of Technology.

References

- Baggaley, W. J., Marsh, S. H., & Close, S. (2007). Interstellar meteors. *Dust in Planetary Systems*.
- Blanchard, J., Lee, N., & Elschot, S. (2022). Meteoroid orbit determination from HPLA radar data. *Icarus*, 386, 115144. <https://doi.org/10.1016/j.icarus.2022.115144>
- Bochkovskiy, A., Wang, C., & Liao, H. (2020). YOLOv4: Optimal speed and accuracy of object detection. arXiv. <https://doi.org/10.48550/arXiv.2004.10934>
- Brown, P., Stober, G., Schult, C., Krzeminski, Z., Cooke, W., & Chau, J. L. (2017). Simultaneous optical and meteor head echo measurements using the Middle Atmosphere Alomar Radar System (MAARSY): Data collection and preliminary analysis. *Planetary and Space Science*, 141, 25–34. <https://doi.org/10.1016/j.pss.2017.04.013>
- Campbell-Brown, M. D., & Close, S. (2007). Meteoroid structure from radar head echoes. *Monthly Notices of the Royal Astronomical Society*, 382(3), 1309–1316. <https://doi.org/10.1111/j.1365-2966.2007.12471.x>
- Chau, J. L., Strelnikova, I., Schult, C., Oppenheim, M. M., Kelley, M. C., Stober, G., & Singer, W. (2014). Nonspecular meteor trails from non-field-aligned irregularities: Can they be explained by presence of charged meteor dust? *Geophysical Research Letters*, 41(10), 3336–3343. <https://doi.org/10.1002/2014GL059922>
- Chau, J. L., & Woodman, R. F. (2004). Observations of meteor-head echoes using the Jicamarca 50 MHz radar in interferometer mode. *Atmospheric Chemistry and Physics*, 4(2), 511–521. <https://doi.org/10.5194/acp-4-511-2004>
- Close, S., Colestock, P., Cox, L., Kelley, M., & Lee, N. (2010). Electromagnetic pulses generated by meteoroid impacts on spacecraft. *Journal of Geophysical Research*, 115(12), 12328. <https://doi.org/10.1029/2010JA015921>
- Close, S., Hamlin, T., Oppenheim, M., Cox, L., & Colestock, P. (2008). Dependence of radar signal strength on frequency and aspect angle of nonspecular meteor trails. *Journal of Geophysical Research*, 113(6), A06203. <https://doi.org/10.1029/2007JA012647>
- Close, S., Hunt, S. M., Minardi, M. J., & McKeen, F. M. (2000). Meteor shower characterization at Kwajalein missile range (Vol. 12). Tech. Rep. No. 1.
- Close, S., Kelley, M., Vertatschitsch, L., Colestock, P., Oppenheim, M., & Yee, J. (2011). Polarization and scattering of a long-duration meteor trail. *Journal of Geophysical Research*, 116(A1), A01309. <https://doi.org/10.1029/2010JA015968>
- Close, S., Oppenheim, M., Durand, D., & Dyrud, L. P. (2005). A new method for determining meteoroid mass from head echo data. *Journal of Geophysical Research*, 110(A9), A09308. <https://doi.org/10.1029/2004JA010950>
- Close, S., Volz, R., Loveland, R., Macdonell, A., Colestock, P., Linscott, I., & Oppenheim, M. (2012). Determining meteoroid bulk densities using a plasma scattering model with high-power large-aperture radar data. *Icarus*, 221(1), 300–309. <https://doi.org/10.1016/j.icarus.2012.07.033>
- Dyrud, L. P., & Janches, D. (2008). Modeling the meteor head echo using Arecibo radar observations. *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(13), 1621–1632. <https://doi.org/10.1016/j.jastp.2008.06.016>
- Dyrud, L. P., Ray, L., Oppenheim, M., Close, S., & Denney, K. (2005). Modelling high-power large-aperture radar meteor trails. *Journal of Atmospheric and Solar-Terrestrial Physics*, 67(13), 1171–1177. <https://doi.org/10.1016/j.jastp.2005.06.016>
- Dyrud, L. P., Wilson, D., Boerve, S., Trulsen, J., Pecseli, H., Close, S., et al. (2008). Plasma and electromagnetic simulations of meteor head echo radar reflections. *Earth, Moon, and Planets*, 102(1–4), 383–394. <https://doi.org/10.1007/s11038-007-9189-8>
- Erickson, P. J., Lind, F. D., Wendelken, S. M., & Faubert, M. A. (2001). Meteor head echo observations using the Millstone Hill UHF incoherent scatter radar system. In *Proceedings of the meteoroids 2001 conference* (pp. 457–463).
- Evans, J. V. (1965). Radio-echo studies of meteors at 68-centimeter wavelength. *Journal of Geophysical Research*, 70(21), 5395–5416. <https://doi.org/10.1029/JZ070i021p05395>
- Guttormsen, G., Fletcher, A. C., & Oppenheim, M. M. (2020). Atomic-scale simulations of meteor ablation. *Journal of Geophysical Research: Space Physics*, 125(9), e2020JA028229. <https://doi.org/10.1029/2020JA028229>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep residual learning for image recognition. arXiv. <https://doi.org/10.48550/arXiv.1512.03385>
- Hedges, T., Lee, N., & Elschot, S. (2022). Meteor head echo analyses from concurrent radar observations at AMISR RResolve Bay, Jicamarca, and Millstone Hill. *Journal of Geophysical Research: Space Physics*, 127(10), e2022JA030709. <https://doi.org/10.1029/2022JA030709>
- Hedges, T., Lee, N., & Elschot, S. (2023). Software, trained models and test data for convolutional neural network trained to detect head echoes in HPLA radar data [Dataset]. *Stanford Digital Repository*. <https://doi.org/10.25740/zn070dk3564>
- Hysell, D. L., Chau, J. L., & Fesen, C. G. (2002). Effects of large horizontal winds on the equatorial electrojet. *Journal of Geophysical Research*, 107(A8), SIA27-1–SIA27-12. <https://doi.org/10.1029/2001JA000217>
- Janches, D., & Chau, J. (2005). Observed diurnal and seasonal behavior of the micrometeor flux using the Arecibo and Jicamarca radars. *Journal of Atmospheric and Solar-Terrestrial Physics*, 67(13), 1196–1210. <https://doi.org/10.1016/j.jastp.2005.06.011>
- Janches, D., Close, S., & Fentzke, J. T. (2008). A comparison of detection sensitivity between ALTAIR and Arecibo meteor observations: Can high power and large aperture radars detect low velocity meteor head-echoes. *Icarus*, 193(1), 105–111. <https://doi.org/10.1016/j.icarus.2007.08.022>
- Janches, D., Dyrud, L. P., Broadley, S. L., & Plane, J. M. (2009). First observation of micrometeoroid differential ablation in the atmosphere. *Geophysical Research Letters*, 36(6), L06101. <https://doi.org/10.1029/2009GL037389>
- Janches, D., Hocking, W., Pifko, S., Hormaechea, J. L., Fritts, D. C., Brunini, C., et al. (2014). Interferometric meteor head echo observations using the Southern Argentina Agile Meteor Radar. *Journal of Geophysical Research: Space Physics*, 119(3), 2269–2287. <https://doi.org/10.1002/2013JA019241>
- Kastinen, D., & Kero, J. (2022). High-altitude meteors detected by the interferometric MU radar. *Monthly Notices of the Royal Astronomical Society*, 517(2), 3024–3033. <https://doi.org/10.1093/mnras/stac2791>
- Kingma, D., & Ba, J. (2017). Adam: A method for stochastic optimization. arXiv. <https://doi.org/10.48550/arXiv.1412.6980>
- Kozlovsky, A., Lukianova, R., & Lester, M. (2020). Occurrence and altitude of the long-lived nonspecular meteor trails during meteor showers at high latitudes. *Journal of Geophysical Research: Space Physics*, 125(8), e2019JA027746. <https://doi.org/10.1029/2019JA027746>
- Li, A., & Close, S. (2015). Mean thermospheric density estimation derived from satellite constellations. *Advances in Space Research*, 56(8), 1645–1657. <https://doi.org/10.1016/j.asr.2015.07.022>
- Li, A., & Close, S. (2016). Neutral density estimation derived from meteoroid measurements using high-power, large-aperture radar. *Journal of Geophysical Research: Atmospheres*, 121(13), 8023–8037. <https://doi.org/10.1002/2015JD024547>
- Li, Y., Galindo, F., Urbina, J., Zhou, Q., & Huang, T. (2022). Meteor detection with a new computer vision approach. *Radio Science*, 57(10), e2022RS007515. <https://doi.org/10.1029/2022RS007515>
- Li, Y., Galindo, F., Urbina, J., Zhou, Q., & Huang, T. (2023). A machine learning algorithm to detect and analyze meteor echoes observed by the jicamarca radar. *Remote Sensing*, 15(16), 4051. <https://doi.org/10.3390/rs15164051>

- Li, Y., Zhou, Q., Scott, M., & Milla, M. (2020). A study on meteor head echo using a probabilistic detection model at Jicamarca. *Journal of Geophysical Research: Space Physics*, 125(1), e2019JA027459. <https://doi.org/10.1029/2019JA027459>
- Limonta, L., Close, S., & Marshall, R. A. (2020). A technique for inferring lower thermospheric neutral density from meteoroid ablation. *Planetary and Space Science*, 180, 104735. <https://doi.org/10.1016/j.pss.2019.104735>
- Marshall, R., Brown, P., & Close, S. (2017). Plasma distributions in meteor head echoes and implications for radar cross section interpretation. *Planetary and Space Science*, 143, 203–208. <https://doi.org/10.1016/j.pss.2016.12.011>
- Matzka, J., Stolle, C., Yamazaki, Y., Bronkalla, O., & Morschhauser, A. (2021). The geomagnetic Kp index and derived indices of geomagnetic activity. *Space Weather*, 19(5), e2020SW002641. <https://doi.org/10.1029/2020SW002641>
- Oppenheim, M. M., & Dimant, Y. S. (2015). First 3-D simulations of meteor plasma dynamics and turbulence. *Geophysical Research Letters*, 42(3), 681–687. <https://doi.org/10.1002/2014GL062411>
- Oppenheim, M. M., Vom Endt, A. F., & Dyrud, L. P. (2000). Electrodynamics of meteor trail evolution in the equatorial E-region ionosphere. *Geophysical Research Letters*, 27(19), 3173–3176. <https://doi.org/10.1029/1999GL000013>
- Panka, P. A., Weryk, R. J., Bruzzone, J., Janches, D., Schult, C., Stober, G., & Hormaechea, J. (2021). An improved method to measure head echoes using a meteor radar. *The Planetary Science Journal*, 2(5), 197. <https://doi.org/10.3847/PSJ/ac22b2>
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). PyTorch: An imperative style, high-performance deep learning library. *arXiv*. <https://doi.org/10.48550/arXiv.1912.01703>
- Popova, O., Sidneva, S., Strelkov, A., & Shuvalov, V. (2001). Formation of disturbed area around fast meteor body. In *Proceedings of the meteoroids 2001 conference* (pp. 237–245).
- Schult, C., Brown, P., Pokorný, P., Stober, G., & Chau, J. L. (2018). A meteoroid stream survey using meteor head echo observations from the middle Atmosphere ALOMAR Radar System (MAARSY). *Icarus*, 309, 177–186. <https://doi.org/10.1016/j.icarus.2018.02.032>
- Schult, C., Stober, G., Chau, J. L., & Latteck, R. (2013). Determination of meteor-head echo trajectories using the interferometric capabilities of maarsy. *Annales Geophysicae*, 31(10), 1843–1851. <https://doi.org/10.5194/angeo-31-1843-2013>
- Sparks, J. J., Janches, D., Nicolls, M., & Heinselman, C. (2009). Seasonal and diurnal variability of the meteor flux at high latitudes observed using PFISR. *Journal of Atmospheric and Solar-Terrestrial Physics*, 71(6–7), 644–652. <https://doi.org/10.1016/j.jastp.2008.08.009>
- Stober, G., Brown, P., Campbell-Brown, M., & Weryk, R. (2021). Triple-frequency meteor radar full wave scattering - Measurements and comparison to theory. *Astronomy and Astrophysics*, 654, A108. <https://doi.org/10.1051/0004-6361/202141470>
- Sugar, G., Marshall, R., Oppenheim, M. M., Dimant, Y. S., & Close, S. (2021). Simulation-derived radar cross sections of a new meteor head plasma distribution model. *Journal of Geophysical Research: Space Physics*, 126(7), e2021JA029171. <https://doi.org/10.1029/2021JA029171>
- Urbina, J., & Briczinski, S. (2011). Transmitter power studies on meteor radar head echo returns. In *2011 XXXth URSI general assembly and scientific symposium*. Blackwell Publishing Ltd. <https://doi.org/10.1029/2009JA014054>
- Volz, R., & Close, S. (2012). Inverse filtering of radar signals using compressed sensing with application to meteors. *Radio Science*, 47(6), RS0N05. <https://doi.org/10.1029/2011RS004889>