WSR-88D Chaff Detection and Characterization using an Optimized Hydrometeor Classification Algorithm

JAMES M. KURDZO, BETTY J. BENNETT, MARK S. VEILLETTE, DAVID J. SMALLEY, EARLE R. WILLIAMS, AND MICHAEL F. DONOVAN

MIT Lincoln Laboratory, Lexington, Massachusetts

ABSTRACT

Chaff presents multiple issues for aviation, air traffic controllers, and the FAA, including false weather identification and areas where flight paths may need to be altered. Chaff is a radar countermeasure commonly released from aircraft across the United States and is comprised of individual metallic strands designed to reflect certain wavelengths. Chaff returns tend to look similar to weather echoes in the reflectivity factor and radial velocity fields, and can appear as clutter, stratiform precipitation, or deep convection to the radar operator or radar algorithms. When polarimetric fields are taken into account, however, discrimination between weather and non-weather echoes has relatively high potential for success. In this work, the operational Hydrometeor Classification Algorithm (HCA) on the WSR-88D is modified to include a chaff class that can be used as input to a Chaff Detection Algorithm (CDA). This new class is designed using human-truthed chaff datasets for the collection and quantification of variable distributions, and the collected chaff cases are leveraged in the tuning of algorithm weights through the use of a metaheuristic optimization. A final CDA uses various image processing techniques to deliver a filtered output. A discussion regarding WSR-88D observations of chaff on a broad scale is provided, with particular attention given to observations of negative differential reflectivity during different stages of chaff fallout. Numerous cases are presented for analysis and characterization, both as an HCA class and as output from the filtered CDA.

1. Introduction

Military chaff has been utilized as a radar countermeasure around the world since World War II (De Martino, 2012). Due to its metallic coating and specific cut lengths to resonate at a given incident frequency of electromagnetic energy (Hessemer 1961; Palermo and Bauer 1965), chaff can generate substantial radar cross sections that can disperse into large “clouds” of distributed targets (Pinson 1975; Harrison and Heinz 1963; Zrnić and Ryzhkov 2004 [referred to as ZR04 from this point forward]; Fig. 1). These targets simulate broad returns to an enemy radar system, leading to the potential for masking aircraft, warships, and missiles from enemy detection (Pode 1960).

The utility and effectiveness of chaff has led to widespread use not only in the battlefield, but, logically, also as a training tool in the United States and abroad. Due to its prevalence in training exercises, chaff can often be found across the national airspace in day-to-day aviation operations. Chaff can be cut to different lengths in order to generate effective returns at a frequency of interest.

Corresponding author address: James M. Kurdzo, Air Traffic Control Systems, MIT Lincoln Laboratory, 244 Wood Street, Lexington, MA 02420. E-mail: James.Kurdzo@ll.mit.edu

This material is based upon work supported by the FEDERAL AVIATION ADMINISTRATION under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the FEDERAL AVIATION ADMINISTRATION.

Since S-band (approximately 10-cm wavelength) radar systems are relatively common across the world, numerous types of chaff are cut to the appropriate half-wavelength size to offer optimal response at S band (Hall et al. 1984; Zrnić and Ryzhkov 1999). Therefore, with 159 S-band WSR-88D radars across the United States and world (Crum and Alberty 1993), chaff is often visible to forecasters, flight controllers, and the general public (e.g., ZR04; Ryzhkov et al. 2005; Murphy et al. 2016).

Chaff can masquerade as stratiform precipitation, convective precipitation, a biological scatterer, combustion debris, and other potential targets owing to similar characteristics, especially in reflectivity factor estimates (Melnikov et al. 2008; Ryzhkov et al. 2005; Murphy et al. 2016). Despite often dispersing in distinct linear segments after release from an aircraft, chaff’s eventual appearance can differ significantly from event to event and can be related to the 3-dimensional wind field, turbulence on multiple scales, and nearby meso- and microscale phenomena. This means that when chaff is released near convective activity, for example, it can be difficult for algorithms, forecasters, and flight controllers to distinguish between chaff and convection using reflectivity factor estimates alone.

Since the deployment of dual-polarimetric capabilities across the WSR-88D fleet was completed in 2013 (Doviak et al. 2000), a number of additional estimates have been made available to radar users (i.e., differential reflectivity, cross-correlation coefficient, and differential phase; Doviak and Zrnić 2006; Bringi and Chandrasekar 2005). These additional estimates support new potential to make identification of chaff more straightforward, leading to better separation of chaff (considered "clutter" to some weather radar users) from targets of interest to the user. However, these additional dual-polarimetric data products are not available to all users (e.g., flight controllers), and can take considerable knowledge and training for accurate application to quickly identify chaff and separate it from weather. While this may be relatively simple for a trained forecaster, other users may find such distinctions to be difficult or impossible to make. For that reason, a viable chaff detector algorithm would be of use to the FAA, weather radar, and National Weather Service (NWS) communities.

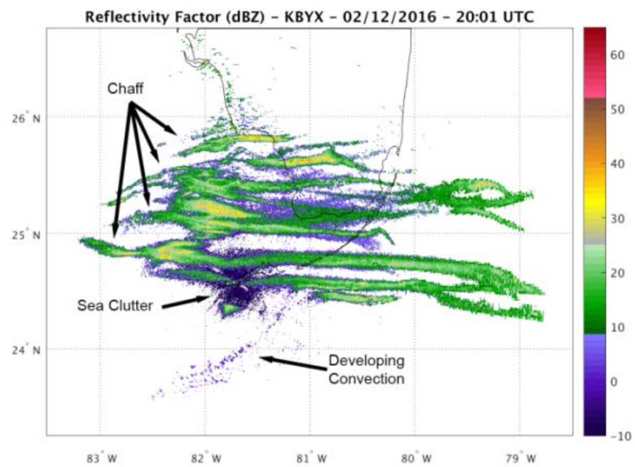

FIG 1: Bands of chaff (horizontal trails) in the reflectivity factor estimate on 12 February 2016 at the KBYX WSR-88D.

Current target classification using the WSR-88D is performed primarily by the Hydrometeor Classification Algorithm (HCA; Park et al. 2009). The HCA utilizes the concept of fuzzy logic to determine the most likely target type based on a series of approximate distributions and weights of different variables collected from radar estimates. In the WSR-88D implementation, Park et al. (2009) describe the distributions and weights for the six following variables:

- Reflectivity Factor
- Differential Reflectivity
- Cross-Correlation Coefficient
- A Linearized Version of Specific Differential Phase
- Standard Deviation of Reflectivity Factor
- Standard Deviation of Differential Phase

Through years of experimentation (e.g., Ryzhkov et al. 2005) and theoretical analysis (e.g., Bringi and Chandrasekar 2005), distributions of the six variables were determined for 10 different target types:

- Ground Clutter/Anomalous Propagation (GC/AP)
- Biological Scatterers (BI)
- Dry Snow (DS)
- Wet Snow (WS)
- Ice Crystals (IC)
- Graupel (GR)
- Big Drops (BD)
- Rain (RA)
- Heavy Rain (HR)
- Rain/Hail Mixture (HA)

These distributions were simplified as non-symmetric trapezoidal functions in either one or two dimensions, convolved with appropriate data quality vectors, related to the altitude of the melting layer, and appropriately weighted to determine the most likely target type for a radar pulse resolution volume.

The HCA method provides a statistical analysis tool to radar users that allows for rapid interpretation of the most likely target by combining much of the available information into a single product. This is especially useful for users who are not as familiar with the intricacies of the dual-polarimetric estimates and desire an easy-to-understand product. The method also provides opportunity for growth into additional target types that have distributions that differentiate them from the existing categories. Essentially there is the opportunity to improve upon classifications that erroneously cloak an alternate scatterer (classification) type. It is this technique that we seek to exploit for the development of a Chaff Detection Algorithm (CDA).

Chaff detection with weather radar has been studied multiple times before, but has revolved solely around single-polarization techniques. Specifically, Kim et al. (2013) focus on feature extraction based on spatial and temporal clustering, while Yu et al. (2016) utilize a tree-initialized genetic algorithm to develop a fuzzy scheme for detecting chaff using single-polarization radar data in South Korea. Chaff has also been studied in the meteorological community as a tracer for storm and cloud entrainment (e.g., Moninger and Kropfli 1987; Jung and Albrecht 2014).

Dual-polarimetric radar is a powerful tool for characterizing chaff that is yet to be fully exploited. The single most important attribute of chaff that allows for its discrimination in dual-polarimetric radar data is its very high aspect ratio as a long, thin conductor. This makes chaff extremely anisotropic, more so than any ice crystal or biological target. For these reasons, it is sensible to move forward with chaff detection in the dual-polarimetric realm.

This paper first describes the data collection strategy, and is followed by a description of methodology, results, and a short discussion regarding observations of negative differential reflectivity in chaff (a new finding of this study compared to ZR04). The paper concludes with a summary of findings and upcoming work on the chaff detection topic.

2. Data Collection

Of key importance to this study is the collection of large datasets in order to provide enough variety to accurately determine the distributions of radar variables in chaff. Previous characterizations of polarimetric variable distributions in chaff with weather radar have been presented (ZR04; Murphy et al.

2016), but with small datasets. In ZR04, for example, only one case at one elevation (0.5°) was surveyed for the distributions shown in their study. In 2004 there were relatively few S-band dual-polarimetric weather radars, but with a large network of dual-polarimetric WSR-88Ds now at the community's disposal, a more extensive database can be developed.

Over the course of nine months, WSR-88D data from 141 chaff cases across the United States were collected. It is important to note that while confident these cases all represent chaff, it cannot be known for sure whether all cases are chaff or a different target. In some cases, combustion debris or wildlife can have similar appearance to chaff, even in the polarimetric estimates. Additionally, the only confirmation of these targets being chaff is from news stories and social media posts by National Weather Service offices (e.g., Fig. 2). Aside from the lack of definitive confirmation of chaff, all of the collected cases have similar characteristics and structure.

FIG 2: Twitter post from NWS San Diego on 28 October 2016 regarding chaff appearing on their WSR-88D.

Despite the existence of chaff in a radar volume, there are often additional targets such as weather, biological scatterers, and/or clutter. Therefore, the chaff returns must be manually separated from the non-chaff returns through the use of human truthing. A human-truthing tool was developed to identify cells of chaff and non-chaff. The tool is shown in Fig. 3, and consists of a simple click-based interface for the user to mark whether a cell or cluster is chaff or non-chaff. A subject matter expert uses all of the available information (i.e., radar variables) to make a decision as to the type of target being observed. After a volume analysis is completed, all of the data associated with chaff and non-chaff are stored in a database for future processing.

Since this method depends on cell clustering, some of the cell edges were not included in the original dataset, so a "lasso" tool implementation was used with data collected in the latter half of the collection period, ensuring all aspects of the chaff clouds were included. In total, 22 cases were used to develop the distributions shown in Section 3, comprising over 500,000 data points.

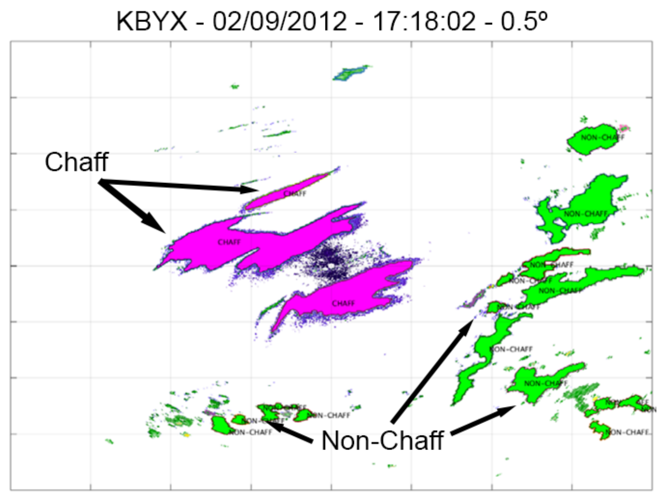

FIG 3: A human-truthing tool for marking chaff and non-chaff was developed for this study at MIT Lincoln Laboratory.

3. Methodology

a. Calculations of distributions

Once a substantial amount of chaff data was human-truthed and stored, distributions of radar variables could be created. These distributions are a critical input metric for the fuzzy logic method utilized in the HCA framework. The distributions of reflectivity factor (Z), differential reflectivity (Z_{DR}), cross-correlation coefficient (ρ_{HV}), and standard deviations of reflectivity factor (SD[Z]) and differential phase (SD[ϕ_{DP}]) are shown in Fig. 4.

Reflectivity factor displays a relatively low signature with a median value of approximately 15 dBZ, while cross-correlation coefficient is represented by a nearly symmetric distribution centered just below ρ_{HV} = 0.6. The standard deviation of reflectivity factor is similar to many other target types with a median value just over 2, while the standard deviation of differential phase is significantly higher than most target types. Since specific differential phase is only calculated for high ρ_{HV} values, that particular variable is not relevant to chaff and is not plotted. It is important to note that these distributions are for all elevations and tilts so as to serve as an input for the current HCA implementation. The distributions, especially Z_{DR}, are different at varying elevations.

Of primary interest is the wide distribution of differential reflectivity, with non-trivial percentages of the distribution ranging across the entire spectrum. This is in stark contrast to multiple previous reports (Zrnić and Ryzhkov 1999; ZR04; Melnikov et al. 2008), specifically ZR04. As stated earlier in the paper, this is due to the sheer volume of data used in this study relative to ZR04, but additional discussion of this observation is provided in Section 5. It is important to note that these distributions were calculated after the Level-II data were passed through a Radar Product Generator (RPG) dual-polarimetric pre-processor simulator, essentially meaning that Level-III data were used to create the distribution plots. While the distributions are slightly different using Level-II data, the trends are similar. This method was chosen due to the fact that the current version of the HCA operates on data that have been processed through the RPG. Note also that the Z_{DR} data have range bounds of +/- 7.9 dB. An extended range Z_{DR} should be available on the WSR-88D in the coming years and it is expected that could possibly enhance the tails of the distribution.

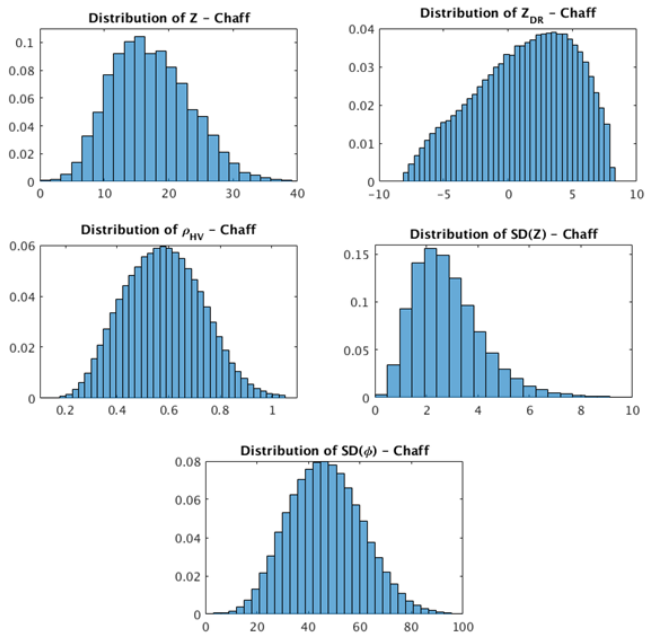

FIG 4: Distributions of Z, Z_{DR}, ρ_{HV}, and standard deviation of Z and ϕ_{DP} in chaff cases across the United States.

b. Optimization of membership functions and weights

The membership functions within the HCA framework are simplified distributions that take the form of a non-symmetric trapezoid. While the distributions shown in Fig. 4 cover a multitude of cases and are therefore expected to be accurate representations of the variable distributions, they must be simplified in a sense that will create a proper chaff class for the HCA. Trapezoids that approximate the distributions are fed into a genetic algorithm optimization framework (Holland 1992; Eiben and Smith 2003) as initial conditions, along with initial weights of 0.5 for each variable. Genetic algorithms represent a set of metaheuristics that step through "generations" of possible solutions made up of bit-sequence combinations of previous generations. Random mutations are applied to avoid local minima in the fitness function, allowing the procedure to operate in a similar nature to the theory of evolution. The four corners of each trapezoid (X_1, X_2, X_3, and X_4) that relate to 0, 1, 1, and 0 probabilities (respectively) are each allowed to deviate slightly from their initializations while the weights are allowed to vary from 0 to 1. Data resolutions are different for each variable, but range from whole integers to two decimal places.

The number of degrees of freedom in the optimization is therefore 30 (not counting any thresholds, covered in Section 3c), with 24 of the degrees relating to trapezoid position and the remaining 6 relating to weights. The weights, *W*, are then convolved with the membership function, *P*, and a quality index vector, *Q*, to form an aggregation value, *A*:

$$A_i = \frac{\sum_{j=1}^{6} W_{ij} Q_j P^i(V_j)}{\sum_{j=1}^{6} W_{ij} Q_j} \tag{1}$$

This summation passes through all six variables (the *j* index) to determine a weighted aggregate score for a given target (the *i* index). The target with the highest aggregation score is the HCA output for a given range gate.

The genetic algorithm utilizes the 30 degrees of freedom as inputs and operates on a series of chaff and non-chaff datasets. The goal is to maximize the number of chaff gates that are reported in the chaff class and minimize the number of non-chaff gates that have their original HCA classification changed. This technique therefore attempts to minimize the effect on existing HCA classes while developing a new chaff class. The fitness function that has yielded the most positive results is a critical success index represented by:

$$f = \frac{S_C}{S_C + P_C} \tag{2}$$

where S_C is the chaff score (a "hit") and P_C is the chaff penalty (a "miss"). The algorithm attempts to minimize *f* by changing the trapezoid positions and weights.

c. Data thresholding

As described in Park et al. (2009), the accuracy of the HCA can be improved by suppressing certain classes with thresholds. These thresholds are purely empirical, meaning that they can be built in to the optimization algorithm to suppress improper classifications of chaff in situations where it is unlikely to be observed. Although multiple thresholds have been experimented with, the primary threshold of interest for chaff has proven to be spectrum width. As shown in Fig. 5, the observed spectrum width in chaff is generally below 4 m s^{-1} and usually below 2 m s^{-1}. As part of the optimization utilized in this paper, a value of 2.9 m s^{-1} was used as a threshold, meaning any gates with a spectrum width higher than 2.9 m s^{-1} are not allowed to be classified as chaff. This particular threshold reduces chaff false alarms in ground clutter and anomalous propagation where spectrum widths are generally higher. The 2.9 m s^{-1} value was determined via optimization as an extra degree of freedom, and is subject to changes in the future. Additionally, this number may change as more cases are included in the optimization. Future work may involve the use of a texture parameter in the spectrum width field.

The use of spectrum width as a data thresholding parameter is based on the assumption that chaff tends to exist in low-turbulence environments. In other words, chaff is generally not found in highly sheared environments. These types of environments would lead to a dramatic widening of the spectrum, and they do occasionally occur. However, the assumption regarding low spectrum width generally performs well, especially when the image processing techniques discussed in the next subsection are applied to fill in any gaps where spectrum width is either elevated or not available.

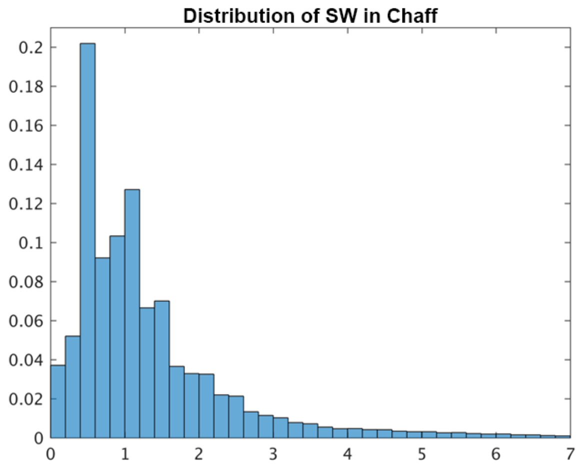

FIG 5: Distribution of spectrum width (SW) in 22 chaff cases across the United States. Units are m s^{-1} along the abscissa and normalized probability along the ordinate.

d. Image processing

While the HCA output provides users with a general idea of which hydrometeors are being detected, it can appear rather noisy at times and is not necessarily ideal if available to flight controllers in this form. These users often prefer to see a single product that can differentiate between two target types (e.g., weather and clutter/chaff). In fact, the preferred practice for FAA is to distill the HCA (and other dual-polarimetric information) into separate products that interpret the situation in terms useful to them (such as icing or hail). Other algorithms exist to differentiate between weather and non-weather (e.g., Krause 2016), but chaff is an important distinction for the FAA. Chaff ingest into aircraft engines can result in more expenses related to maintenance and cleaning and it is therefore preferred to avoid flying into chaff clouds. For this reason, a binary product that determines whether a target is chaff or not chaff may potentially be ideal for the FAA and flight controllers.

Since the raw HCA output can contain noisy data and it is desired to provide binary chaff information in a cellular sense (i.e., an entire chaff "cloud" should be marked as chaff rather than only parts of it), some form of image processing must be applied to the HCA chaff class. It is also critical to FAA needs to avoid labeling weather as chaff, so the image processing aides in mitigating this. The following steps are taken during the image processing phase of the algorithm:

1. Calculate HCA output
2. Separate chaff and non-chaff
3. Median filter each target type (chaff and non-chaff)
4. Dilate and close each target type
5. Filter out the non-chaff classes from the chaff class
6. Final dilation and closing

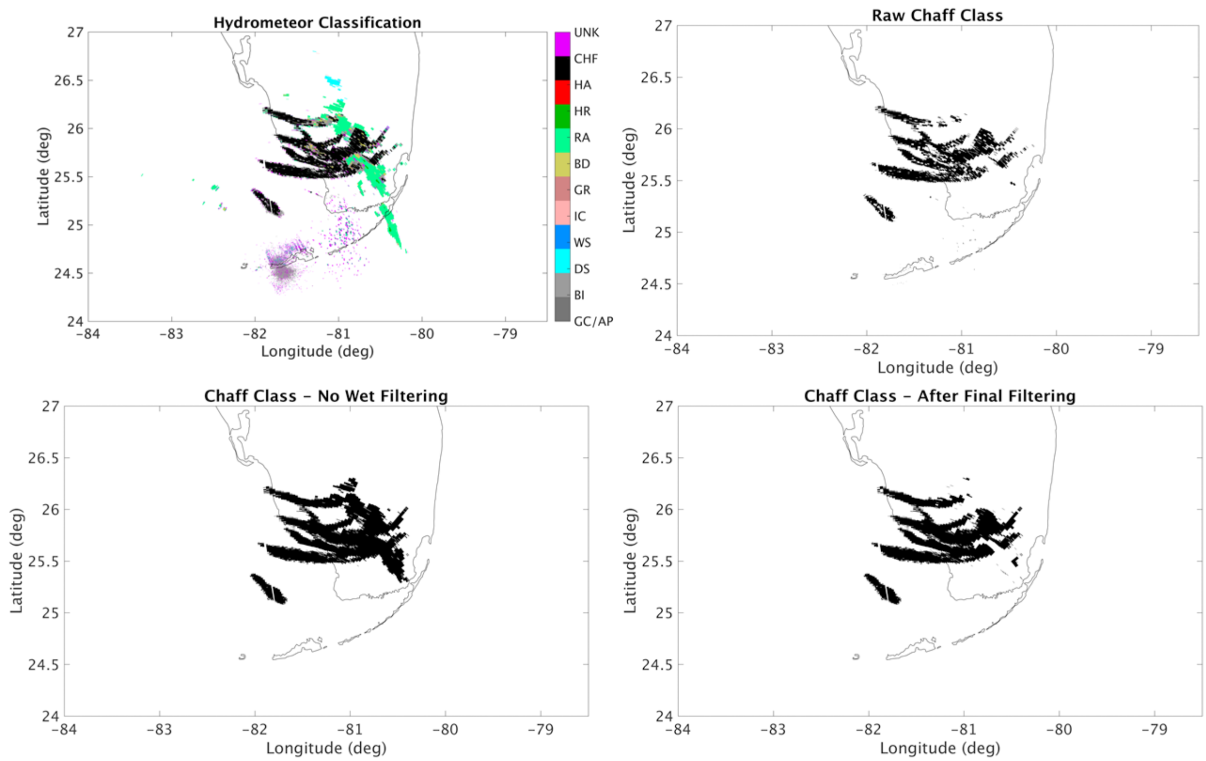

FIG 6: Illustration of four of the six image processing steps for a mixed chaff/weather case at the KBYX WSR-88D site on 23 February, 2016. Step 1 (top left) is the HCA output; Step 2 (top right) is the separation of chaff from non-chaff; Step 4 (bottom left) is the result of dilation/closing of the chaff class; Step 6 (bottom right) is the final chaff product after accounting for the weather targets.

The core aspects of this technique revolve around binary morphological operators such as median filtering, dilation, and closing. The median filtering cleans up initial noise in areas of small detections, while the dilation and closing fill in cells that weren't fully detected as chaff. The filters can be rather large spatially depending on the type of target that is detected and the resulting noise levels, but these specifics are still being developed.

While the chaff class is filtered in this way, the non-chaff targets are also median filtered, dilated, and closed. Non-chaff targets are considered to be any liquid targets below the melting layer other than the HCA big drops classification and any ice-phase targets that meet appropriate SD(ϕ_{DP}) and ϕ_{DP} requirements. This result is convolved with the filtered chaff targets to remove any possibility of weather targets being labeled as chaff. Due to the dilation and closing of the non-weather target type, a small buffer is created around the weather as a conservative measure. A final dilation and closing is performed on the chaff class, taking into account the filtered weather targets. An illustration of steps 1, 2, 4, and the final result (step 6) is shown in Fig. 6.

e. Machine learning

While exploring the topic of chaff detection and determining the appropriate membership functions and weights for a chaff class in the existing HCA, it was found that a number of errors arose in cases of sea clutter. After closer analysis, it was discovered that the radar variable distributions for chaff and sea clutter are relatively similar to each other. Specifically, ρ_{HV} values are low due to the lack of homogeneity and ϕ_{DP} values are high and variable, leading to a high texture parameter. While this is an ongoing problem being worked on, it is important to note that initial attempts to correct this issue with machine learning have been successful. Although beyond the scope of this paper, a short description of this final step is warranted.

The resulting chaff "clusters" (a cell of chaff, similar to thunderstorm cell detection) represent regional statistics that provide distributions of variables within an entire cell rather than at one range gate (local statistics). These characteristics show some differences between sea clutter and chaff and can be "taught" to a machine learning algorithm. A support vector machine (SVM) approach (Bishop 2007) is currently being adapted to look into this as a possible solution. An SVM is a supervised learning technique that is fed by large, truthed training datasets. By finding patterns in data from the training datasets, it can predict the appropriate classification of new datasets. As long as a large enough dataset is used for training, no re-training or unsupervised learning is necessary during processing of new data. Initial results have been promising but need additional work in the near future.

4. Results

Three chaff cases are presented in this section as representative examples for the chaff detector's performance. Each case includes a six-panel plot of Z, Z_{DR}, ρ_{HV}, ϕ_{DP}, the HCA output with the chaff class, and the final CDA output. The HCA output also includes a class for sea clutter (see section 3e), although this is a work in progress. None of the cases shown contain significant sea clutter that would be detected by the algorithm. Additionally, a class marked as "NL" (non-living) demarcates areas where a chaff-like target was detected but neither chaff or sea clutter could be determined as a highly likely target. More details on this can be found in the discussion section. While the chosen cases are only a subset of the cases collected, they show the general performance capabilities of the algorithm. None of the cases presented were confirmed by social media posts from the NWS or other sources, but they follow the same characteristics as cases that were cited by the NWS.

a. Case 1: 26 January, 2016 – KLTX (Wilmington, NC)

A series of chaff releases at different heights and times resulted in an east-west oriented chaff cloud over the ocean, a chaff cloud over extreme southeastern North Carolina, and a number of small areas of chaff in east-west orientations extending in a line from north to south. During the release, a line of convective showers formed in a north-south orientation near the middle of the large chaff cloud off the coast. A sample of Z, Z_{DR}, ρ_{HV}, ϕ_{DP}, HCA, and CDA outputs is shown in Fig. 7 for the base elevation scan.

The chaff is particularly simple to see in the ϕ_{DP} field in this case due to the lack of significant ground clutter. The areas of chaff are generally marked by regions of high ϕ_{DP}, which is in stark contrast to the areas of relatively low and constant ϕ_{DP} in weather. ϕ_{DP} alone can be a difficult tool to use for determining whether a target is chaff because of different calibrations at different radar sites as well as the existence of higher ϕ_{DP} values in heavy rain, hail, etc. For these reasons, the texture parameter of ϕ_{DP}, or the standard deviation of ϕ_{DP}, is a heavily weighted tool for distinguishing between chaff and other targets. This case shows the high SD(ϕ_{DP}) in the chaff clouds consistently.

The HCA returns a chaff classification for the majority of the areas of chaff, although some areas of the large east-west oriented band are marked by ground clutter and big drops. The image processing that results in the final CDA output fills in these gaps to produce single-cell chaff cloud representations. Also of note in this case is the lack of weather censoring as the convective line crosses the chaff cloud over the ocean. This is because the radar resolution volumes are filled predominantly with chaff, leading to the correct designation being assumed to be chaff. This is a common problem in mixed chaff and weather cases, but when the returns are dominated by weather, the range gate is protected from being labeled as chaff.

Z_{DR} is distributed non-uniformly at the base elevation scan, with strongly positive values to the east (higher elevations) and strongly negative values to the west (closer to the radar) in the large east-west oriented cell. The main chaff cell over land contains areas of high and low Z_{DR}, while some of the smaller cells consist mostly of positive Z_{DR} values and some consist of predominantly negative Z_{DR} values. This case generally shows

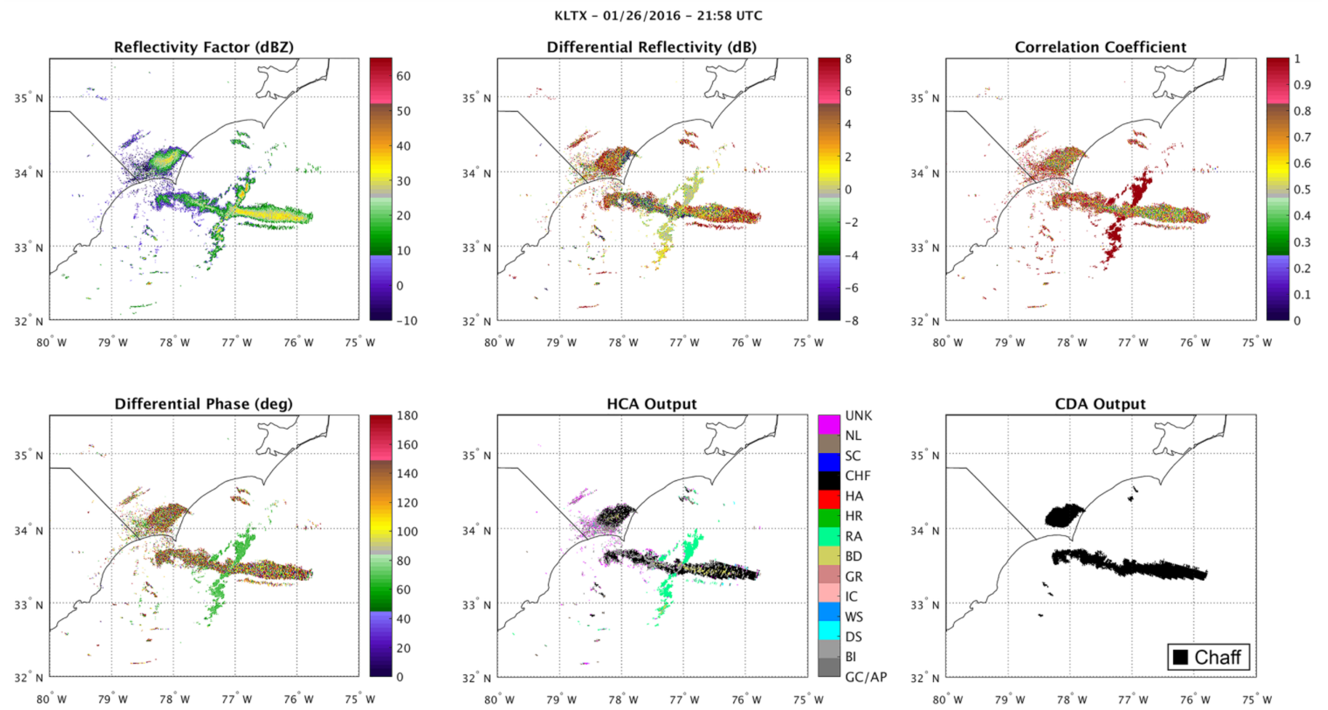

FIG 7: From left to right, top to bottom: Z, Z_{DR}, ρ_{HV}, ϕ_{DP}, HCA output, and CDA output from a mixed chaff and weather case on 26 January 2016 at 2158 UTC at the 0.5° elevation angle from KLTX. The CDA properly identifies the two largest chaff cells but does not detect all of the smaller chaff cells to the south.

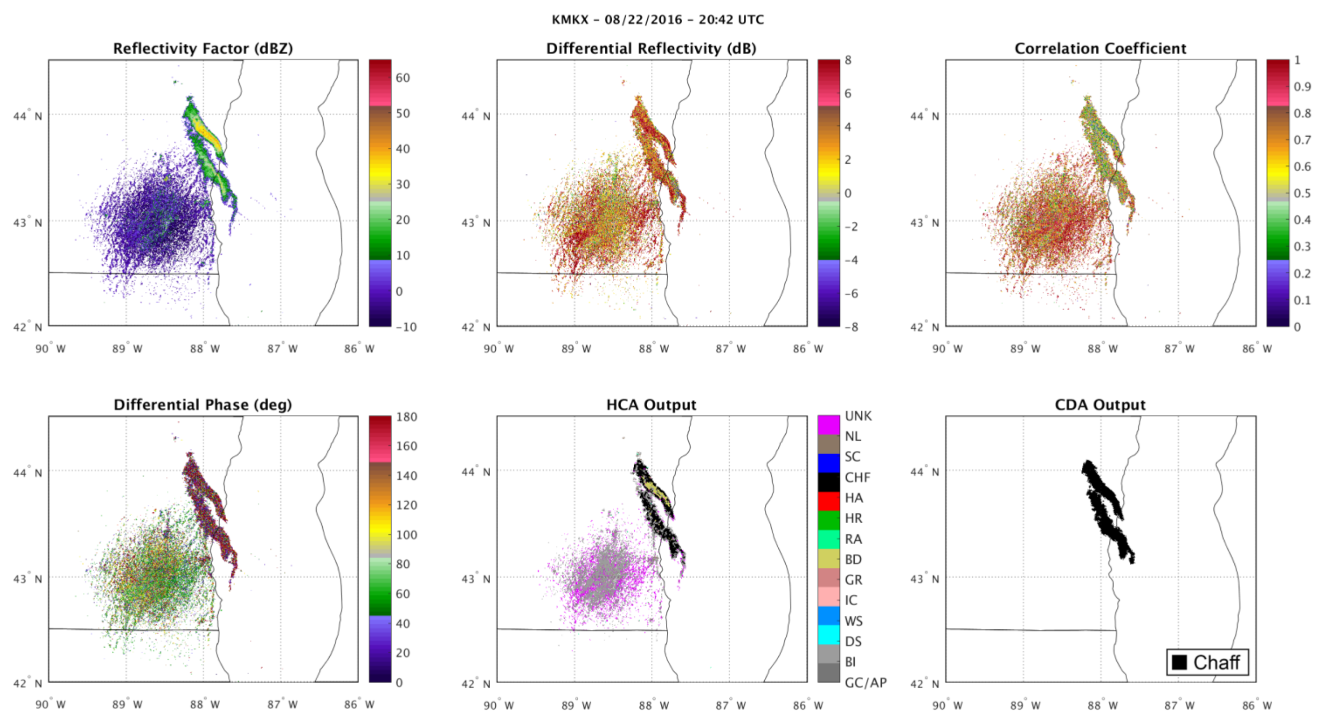

FIG 8: From left to right, top to bottom: Z, Z_{DR}, ρ_{HV}, ϕ_{DP}, HCA output, and CDA output from a mixed chaff and clutter case on 22 August 2016 at 2042 UTC at the 0.5° elevation angle from KMKX. The CDA properly identifies the two chaff cells and does not detect any chaff in the clutter near the radar.

negative Z_{DR} at lower elevations and positive Z_{DR} at higher elevations, but some areas do not follow the overall trend. Some cells seem to be influenced by nearby weather or are at different fallout stages, possibly indicating that not enough time has passed for the chaff to become vertically oriented.

b. Case 2: 22 August, 2016 – KMKX (Milwaukee, WI)

Two separate but closely spaced chaff releases occurred to the northeast of Milwaukee, WI. Mid-level winds from the southeast caused the chaff to fall out towards central WI, resulting in a 150-km band of chaff to appear on the 0.5° tilt of KMKX at 2042 UTC (Fig. 8). This case demonstrates the difficulty of using just ϕ_{DP}, for example, to differentiate chaff from other targets. Despite a high standard deviation, the values range substantially from nearly 0° to 360°, indicating not only a higher standard deviation, but the existence of phase wrapping that can be difficult to quantify in noisy fields.

The chaff clouds, oriented roughly north-south, overlap the outer edges of the clutter. While many areas of the clutter exhibit less texture in ϕ_{DP}, some areas look similar to the chaff. In this case, the other parameters are combined to differentiate the chaff from the non-chaff targets. Although not shown, the spectrum width in the clutter is generally elevated relative to the chaff, causing fewer classifications of chaff in the HCA. The image processing smooths over these small detections, eliminating them from the final CDA.

Also of note is the existence of big drop detections in chaff clouds, indicating similarities between the membership functions of big drops and chaff. One possible solution to this is to limit the big drop class to maximum ϕ_{DP} texture parameters to avoid big drops from appearing from chaff. However, this could potentially cause a cascading effect in the HCA that would need to be properly validated.

It is worth noting that big drops consist of primarily positive Z_{DR} values due to their oblate spheroidal shape. The overlap between the big drop and chaff classes is therefore an interesting observation, but given the wide Z_{DR} range of chaff, one would expect that big drop classifications may appear only in the areas of strongly positive Z_{DR}, which is supported in this case. Z_{DR} is generally positive in this case, even at the base elevation scan. This was towards the tail end of the chaff cloud fallout, a period that tends to display positive Z_{DR} values even at lower elevations. These observations draw into question the discussion in Section 5 regarding the fair-weather electric field, but may also point to additional mechanisms that are yet to be explored for chaff orientation, particularly in the lower troposphere.

c. Case 3: 1 November, 2016 – KSFX (Pocatello, ID)

The final example is a mixed chaff, weather, and clutter case from a chaff release west of Pocatello, ID. Northeasterly winds resulted in the appearance of two chaff clouds approximately 150-200 km southwest of the radar site on the 0.5° tilt. These clouds can be verified (as best able) by tracking their occurrences at higher elevations over time. Both exhibit similar characteristics to the previous two cases, with generally high ϕ_{DP} and a relatively high texture parameter for ϕ_{DP}. However, the chaff cloud to the east appears to be mixing with rain or snow showers and is therefore influenced by smaller overall ϕ_{DP} values and higher ρ_{HV}. For this reason, the HCA classifies a series of "unknown" gates since no clear distinction can be made among any of the classes. After the image processing algorithm is completed, these cells are removed and only the western-most chaff cloud is detected in the CDA.

The HCA output in the western cloud is almost uniformly chaff, lacking the big drop classifications from the previous two cases. Different types of chaff (wavelengths, materials, etc.), varying calibrations between radars, and different atmospheric conditions may vastly alter the original determination of target types in the HCA. However, as shown across the three cases, the image processing results in a consistent end product.

Z_{DR} in this case is mostly negative, although the outer edges of the cell display positive Z_{DR} values. Z_{DR} values at higher elevations were generally positive (not shown). This case suggests that some aspect of size sorting may be ongoing, although it is not clear as to what the exact sorting mechanism could be, or its scale.

5. Discussion

a. CDA performance and ongoing machine learning work

Performance of the HCA chaff classification and the CDA output have generally met expectations, although there are certain cases that need additional work. First, some cases in clutter can generate clusters of chaff classifications that are large enough that the image processing cannot remove them. This issue is exacerbated by the fact that range gate cells are elongated close to the radar compared to farther distances where they are more symmetric. Since clutter is most prominent in these closer cells, this can cause false positives. This is a possible area for machine learning to aide in removing these detections.

Second, as mentioned previously, sea clutter has a significant overlap of membership functions with chaff, specifically in the critical areas of ϕ_{DP} texture and low spectrum width. For this reason, the regional statistics afforded by utilizing cell clustering and a machine learning approach can be helpful for determining whether cells are sea clutter or chaff. The biggest issue with this technique is that occasionally, chaff cells can be "connected" to sea clutter cells, making the cells cluster together as one. The machine learning algorithm then must make a decision as to which classification to mark the cell. This could potentially be mitigated through a linear time-invariant filter that slides backwards in time to keep continuity between cells. An approach similar to this could also prevent small chaff cells from oscillating between detections and non-detections.

Finally, the use of an overarching class such as the "NL" designation in the previous examples can help label targets that are not chaff or sea clutter but happen to fall into the larger bin with similar characteristics. An example of such a target could be

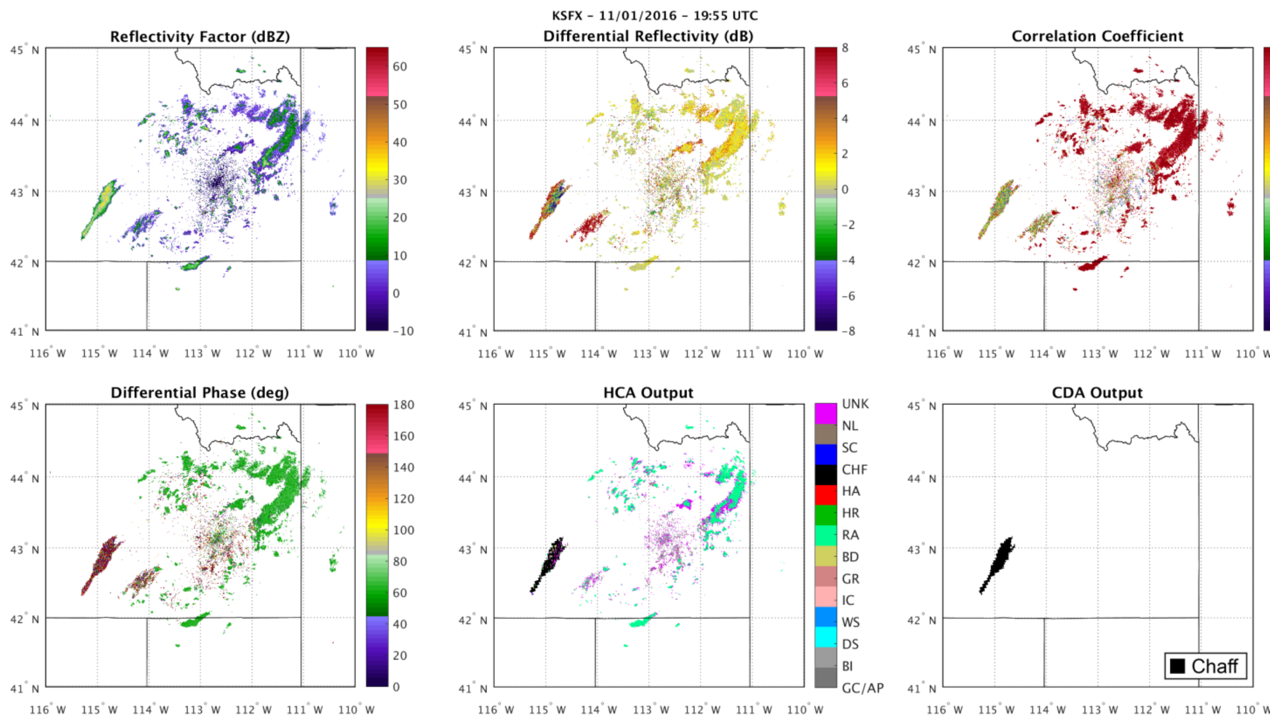

FIG 9: From left to right, top to bottom: Z, Z_{DR}, ρ_{HV}, ϕ_{DP}, HCA output, and CDA output from a mixed chaff, weather, and clutter case on 1 November 2016 at 1955 UTC at the 0.5° elevation angle from KSFX. The CDA properly identifies the chaff cell to the southwest of the radar and does not detect any chaff in the areas of weather or the clutter near the radar.

smoke (specifically ash or combustion debris) associated with a forest fire; these targets take on vastly different characteristics based on what is being burned, but can often fall into similar categories as chaff and sea clutter (Melnikov et al. 2008). The machine learning approach therefore takes anything classified as non-living and puts it in a bin of chaff, sea clutter, or NL. This information could be fed back into the HCA results for greater accuracy. This technique results in small NL detections in the HCA in areas of clutter, but in practice, these results have not deviated significantly in variability compared with the current HCA.

b. Negative Z_{DR} in chaff

While investigating distributions of radar variables in chaff, it was discovered that a non-trivial number of chaff clouds contain strong negative Z_{DR} signals, a new finding from that reported in ZR04. As shown in Fig. 4, the distribution of Z_{DR} is certainly not Gaussian (i.e., there is a skew towards positive Z_{DR} values), but there is a relatively strong signal on the negative end of the Z_{DR} scale that was not seen in ZR04 or other reports. Interestingly, a dependence on height has been seen in the data, leading to the desire for further investigation.

An example of the distribution of Z_{DR} with height for a case at the Key West, FL KBYX WSR-88D site on 17 June 2016 is shown in Fig. 10. A positive slope of approximately 0.29 km dB^{-1} is evident, with generally negative Z_{DR} below 7 km and generally positive Z_{DR} above. Although this is only one case spread over the span of 3 hours, this signature is seen in virtually *all* of the chaff cases analyzed: negative Z_{DR} at the bottom of the chaff cloud transitioning to positive Z_{DR} at the top of the chaff cloud. The exact elevations above ground level deviate somewhat from case to case, especially at different elevations above sea level, but the trend is clear across hundreds of cases.

Two different hypotheses come to mind as to why these observations may be occurring. The most likely scenario is that the fair-weather electric field (Chalmers 1967), which is stronger at lower elevations, is orienting the chaff particles vertically at lower heights. Initial calculations comparing results from Weinheimer and Few (1987) with fall speed estimates taken from Jiusto and Eadie (1963) indicate that the order of magnitude of the fair-weather electric field combined with the aspect ratio of chaff and its effective fall speed could result in vertically dominated orientations, especially at lower elevations. Future work using physical chaff and a controlled electric field could validate this theory and is part of upcoming plans.

Additionally, chaff has been known to clump (Arnott et al. 2004; ZR04), resulting in the need for anti-clumping agents when chaff is released. The authors can verify from experimentation with physical chaff that clumping is a common occurrence. It is possible that clumped chaff falls in a different manner than non-clumped chaff, resulting in a different Z_{DR} signature. Individual chaff strands generally fall in a horizontal fashion when a strong, vertically oriented electric field is not applied, but it is unknown how clumping would affect these ratios.

It would make sense that clumped chaff would be denser and produce less aerodynamic torque and therefore fall out faster than non-clumped chaff. If these clumps were to orient in a more vertical fashion, negative Z_{DR} at lower elevations could be realized. It is suspected that the final Z_{DR} signatures are influenced by both clumping and the electric field. An important area for upcoming work will be analyzing Z_{DR} distributions in cases near convective activity where the electric field is stronger (and reversed) to see if there is an appreciable difference compared with fair-weather cases.

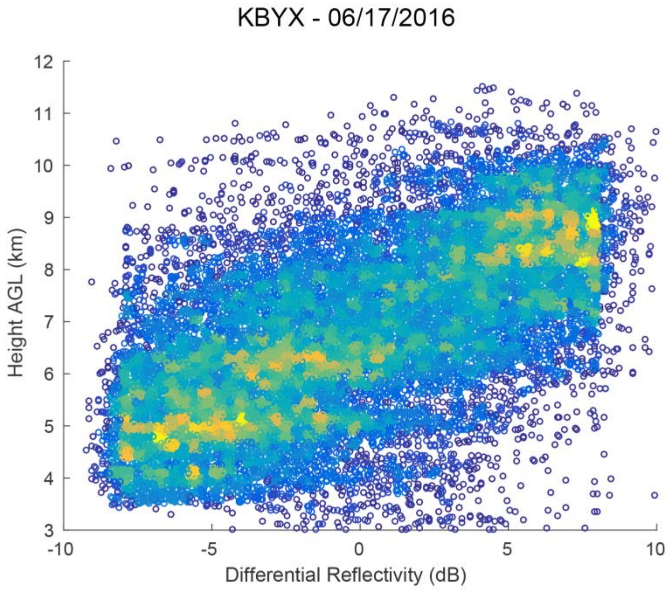

FIG 10: Differential reflectivity vs. height for a three-hour span during a chaff case at KBYX on 17 June 2016. Note the dependence on height, with negative Z_{DR} more likely at lower elevations. Colors indicate relative concentration of values, with hotter (yellow) colors showing higher concentrations.

6. Conclusions and Upcoming Work

This paper has presented preliminary findings regarding the performance of a fuzzy logic-based chaff detection algorithm, or CDA. The CDA consists of an additional classification category added to the existing WSR-88D HCA, combined with an image processing algorithm that produces a binary yes/no chaff product intended as information for FAA flight controllers and other users. It has been found that chaff has similar distributions of radar variables to sea clutter, therefore necessitating further examination of machine learning techniques that utilize regional statistics to separate chaff echoes from sea clutter echoes. These initial results indicate that the use of a “non-living” sub-class in the HCA that can be separated into chaff and sea clutter returns has potential applications in future WSR-88D research and operations. Additionally, repeatable observations of negative Z_{DR} in chaff clouds contrast with previously published reports, resulting in the desire to further investigate the cause and commonality of these occurrences.

Upcoming work includes fine-tuning of the HCA membership functions and weights with larger datasets using the previously developed genetic algorithm method. This tuning should particularly take into account sea clutter targets, which were not considered in this initial implementation. In addition, larger datasets for the SVM training are required, as are more accurate regional statistical measures that will improve performance of the machine learning algorithm. After these steps are complete, implementation in the WSR-88D RPG will allow for large-scale analysis of performance. Additional work regarding the observations of negative Z_{DR} in chaff is also expected in order to better understand the strengths and weaknesses of the algorithm. Considerations for an expected expanded Z_{DR} range will eventually also be addressed.

Acknowledgements. Discussions with Ashish Banerjee, John Cho, Gabriel Elkin, Elaine Griffin, Alexander Ryzhkov, Marilyn Wolfson, and Dusan Zrnić contributed immensely to this work.

REFERENCES

Arnott, W. P. and coauthors, 2004: *Determination of Radar Chaff Diameter Distributon, Fall Speed, and Concentration in the Atmosphere by use of the NEXRAD Radar*. Desert Research Institute.

Bishop, C., 2007: *Pattern Recognition and Machine Learning*. Springer.

Bringi, V. N. and V. Chandrasekar, 2005: *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge University Press.

Chalmers, J. A., 1967: *Atmospheric Electricity*. Pergamon Press.

Crum, T. and R. L. Alberty, 1993: The WSR-88D and the WSR-88D operational support facility. *Bull. of the Amer. Meteor. Soc.*, **74**, 1669-1687.

De Martino, A., 2012: *Introduction to Modern EW Systems*. Artech House.

Doviak, R. J. and coauthors, 2000: Considerations for polarimetric upgrades to operational WSR-88D radars. *J. Atmos. Oceanic Technol.*, **17**, 257-278.

Doviak, R. J. and D. S. Zrnić, 2006: *Doppler Radar and Weather Observations*. 2nd ed., Dover.

Eiben, A. and J. Smith, 2003: *Introduction to Evolutionary Computing*. Springer.

Hall, P. M., J. W. F. Goddard, and S. M. Cherry, 1984: Identification of hydrometeors and other targets by dual-polarization radar. *Radio Science*, **19**, 132-140.

Harrison, C. and R. Heinz, 1963: On the radar cross section of rods, tubes, and strips of finite conductivity. *IEEE Trans. on Antennas and Propagation*, **11**, 459-468.

Hessemer, R., 1961: Scatter communications with radar chaff. *IRE Trans. on Antennas and Propagation*, **9**, 211-217.

Holland, J., 1992: Adaptation in natural and artificial systems. MIT Press.

Jiusto, J. E. and W. J. Eadie, 1963: Terminal fall velocity of chaff. *J. Geophys. Res.*, **68**, 2858-2861.

Jung, E. and B. Albrecht, 2014: Use of radar chaff for studying circulations in and around shallow cumulus clouds. *J. Appl. Meteor. Climatol.*, **53**, 2058-2071.

Kim, Y. and coauthors, 2013: Real-time detection and filtering of chaff clutter from single-polarization Doppler radar data. *J. Atmos. Oceanic Technol.*, **30**, 873-895.

Krause, J. M., 2016: A simple algorithm to discriminate between meteorological and nonmeteorological radar echoes. *J. Atmos. and Oceanic Technol.*, **33**, 1875-1885.

Melnikov, V. M., D. S. Zrnić, R. M. Rabin, and P. Zhang, 2008: Radar polarimetric dignatures of fire plumes in Oklahoma. *Geophysical Research Letters*, **35**, L14815.

Moninger, W. R. and R. A. Kropfli, 1987: A technique to measure entrainment in cloud by dual-polarization radar and chaff. *J. Atmos. Oceanic Technol.*, **4**, 75-83.

Murphy, T. A., R. A. Wade, and B. C. Carcione, 2016: Observations and operational considerations of the 4 June 2013 chaff event in northern Alabama. *J. Operational Meteor.*, **4**.

Palermo, C. J. and L. H. Bauer, 1965: Bistatic scattering cross section of chaff dipoles with application to communications. *Proc. of the IEEE,* **53**, 1119-1121.

Park, H. A., A. Ryzhkov, D. Zrnić, and K. Kim, 2009: The hydrometeor classification algorithm for the polarimetric WSR-88D: Description and application to an MCS. *Wea. Forecasting*, **24**, 730-748.

Pinson, C. C., 1975: Chaff cloud signature. *IEEE International Symposium on Electromagnetic Compatibility*, San Antonio, TX.

Pode, L., 1960: An approach to the study of automatic target tracking. *Trans. of the AIEE, Part I: Communication and Electronics*, **78**, 806-811.

Ryzhkov, A. V. and coauthors, 2005: The joint polarization experiment: Polarimetric rainfall measurements and hydrometeor classification. *Bull. Amer. Meteor. Soc.*, **86**, 809-824.

Weinheimer, A. J. and A. A. Few, 1987: The electric field alignment of ice particles in thunderstorms. *J. Geophys. Res.*, **92**, 14833-14844.

Yu, J., H. Lee, Y. Jeong, and S. Kim, 2016: Identifying chaff echoes in weather radar data using tree-initialized fuzzy rule-based classifier. *IEEE International Conference on Fuzzy Systems*, Vancouver, Canada.

Zrnić, D. and A. Ryzhkov, 2004: Polarimetric properties of chaff. *J. Atmos. Oceanic. Technol.*, **21**, 1017-1024.

Zrnić, D. S. and A. V. Ryzhkov, 1999: Polarimetry for weather surveillance radars. *Bull. of the Amer. Meteor. Soc.*, **80**, 389-406.