

Project Report
ATC-454

Radio Frequency Interference Censoring Scheme for Canadian Weather Radar

J. Y. N. Cho

12 July 2023

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



This report is the result of studies performed at Lincoln Laboratory, a federally funded research and development center operated by Massachusetts Institute of Technology.

© 2023 Massachusetts Institute of Technology

Massachusetts Institute of Technology
Lincoln Laboratory

Radio Frequency Interference Censoring Scheme for
Canadian Weather Radar

J. Y. N. Cho
Group 43

Project Report ATC-454
12 July 2023

Lexington

Massachusetts

This page intentionally left blank.

ABSTRACT

An automated scheme is developed for the upgraded S-band polarimetric Canadian weather radars to detect and censor radio frequency interference from wireless communication devices. The suite of algorithms employed in this scheme effectively identifies and edits out interference-contaminated reflectivity data, while preserving data dominated by weather signals. This scheme was implemented in the NextGen Weather Processor test reference system for continuous real-time testing, and is expected to be incorporated into the new Canadian Aviation Weather Systems.

This page intentionally left blank.

ACKNOWLEDGMENTS

I would like to thank Daniel Michelson (Environment and Climate Change Canada) for providing needed information about the upgraded Canadian weather radars. Thanks also to Danielle Morse (MIT Lincoln Laboratory) for collecting relevant radar data and Evelyn Mann (MIT Lincoln Laboratory) for implementing the algorithms on NWP TRS for real-time testing.

This page intentionally left blank.

TABLE OF CONTENTS

	Page
Abstract	i
Acknowledgments	iii
List of Illustrations	vii
List of Tables	xi
1. INTRODUCTION	1
2. RFI CHARACTERISTICS	3
2.1 Visual Examination	3
2.2 Statistical Analysis	8
3. RFI DETECTION AND CENSORING ALGORITHM	15
3.1 Algorithm Framework	15
3.2 Polarimetric RFI Censor	16
3.3 Non-Polarimetric Spike Censor	18
3.4 Non-Polarimetric Speckle Censor	23
4. EXAMPLE RESULTS	25
5. CONCLUSIONS	29
Glossary	31
References	33

This page intentionally left blank.

LIST OF ILLUSTRATIONS

Figure No.		Page
2-1.	Range-azimuth reflectivity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	4
2-2.	Range-azimuth uncorrected reflectivity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	4
2-3.	Range-azimuth radial velocity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	5
2-4.	Range-azimuth Doppler spectrum width plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	5
2-5.	Range-azimuth differential reflectivity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	5
2-6.	Range-azimuth co-polar correlation coefficient plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	6
2-7.	Range-azimuth differential phase plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	6
2-8.	Range-azimuth uncorrected differential phase plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	6
2-9.	Range-azimuth specific differential phase plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	7
2-10.	Range-azimuth signal quality index plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	7
2-11.	Range-azimuth plot of co-polar correlation coefficient along-range variance from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	8

LIST OF ILLUSTRATIONS (Continued)

Figure No.		Page
2-12.	Range-azimuth plot of uncorrected differential phase along-range circular variance from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .	8
2-13.	Range-azimuth plot of horizontal reflectivity from the Britt, Ontario weather radar at 22:05:15 UT on 25 January 2023, at an elevation angle of 0.4° .	9
2-14.	Range-azimuth plot of horizontal reflectivity from the Landrienne, Quebec weather radar at 00:41:05 UT on 12 November 2022, at an elevation angle of 0.3° .	9
2-15.	Range-azimuth plot of horizontal reflectivity from the Franktown, Ontario weather radar at 14:58:42 UT on 14 February 2023, at an elevation angle of 0.8° .	9
2-16.	Range-azimuth plot of horizontal reflectivity from the Franktown, Ontario weather radar at 14:59:16 UT on 14 February 2023, at an elevation angle of 0.4° .	10
2-17.	Histograms of co-polar correlation coefficient for the weather (left) and RFI (right) case data sets. The bin width is 0.02.	10
2-18.	Histograms of SQI for the weather (left) and RFI (right) case data sets. The bin width is 0.02.	11
2-19.	Histograms of co-polar correlation coefficient along-range variance for the weather (left) and RFI (right) case data sets. The bin width is 0.0003.	11
2-20.	Histograms of uncorrected differential phase along-range circular variance for the weather (left) and RFI (right) case data sets. The bin width is 0.02.	12
2-21.	ROC curves for RFI detection against weather for co-polar correlation coefficient, SQI, co-polar correlation coefficient along-range variance, and uncorrected differential phase along-range circular variance. The left plot shows the full domain, while the right plot zooms in on the upper lefthand corner.	13
3-1.	RFI detection and censoring flow diagram.	15
3-2.	Illustrative examples of spike detection using the recommended parameter settings.	22

LIST OF ILLUSTRATIONS (Continued)

Figure No.		Page
4-1.	Range-azimuth plot of horizontal reflectivity from the Franktown, Ontario weather radar at 23:59:15 UT on 11 November 2022, at an elevation angle of 0.4° , without RFI censoring (top), with polarimetric RFI censoring (second from top), with polarimetric and spike censoring (second from bottom), and with polarimetric, spike, and speckle censoring (bottom).	25
4-2.	Range-azimuth plot of horizontal reflectivity from the Cold Lake, Alberta weather radar at 21:59:05 UT on 25 January 2023, at an elevation angle of 0.3° , without RFI censoring (top), with polarimetric RFI censoring (second from top), with polarimetric and spike censoring (second from bottom), and with polarimetric, spike, and speckle censoring (bottom).	26
4-3.	Range-azimuth plot of horizontal reflectivity from the King City, Ontario weather radar at 00:11:16 UT on 12 November 2022, at an elevation angle of 0.4° , without RFI censoring (top), with K_{dp} invalid data masking (middle), and with RFI censoring (bottom).	27
4-4.	Range-azimuth plot of horizontal reflectivity from the Landrienne, Quebec weather radar at 23:59:05 UT on 11 November 2022, at an elevation angle of 0.3° , without RFI censoring (top), with K_{dp} invalid data masking (middle), and with RFI censoring (bottom).	28

This page intentionally left blank.

LIST OF TABLES

Table No.		Page
2-1	Optimal RFI Detection Thresholds	13

This page intentionally left blank.

1. INTRODUCTION

Radio frequency interference (RFI), especially from wireless devices (WDs) used for communication, is a significant contamination source for weather radar data (Saltikoff 2016). Although most prevalent as direct in-band interference at C band, radars operating in S band can also suffer from indirect out-of-band WD RFI. Radars themselves can also interfere with each other, but due to strict channel assignments and typically low transmission duty cycles, RFI caused by radar is less common and more easily filtered by front-end interference filters when it does occur (Cho 2017).

The Canadian national weather radar network is in the process of being upgraded from single-polarization C-band systems to dual-polarization S-band systems manufactured by Leonardo (ECCC 2023). Somewhat surprisingly, the new radars have been plagued by RFI at many sites (more so than with the old C-band radars), and the cause appears to be out-of-band interference from WDs. Investigation is ongoing to pinpoint the cause at each site and to develop a long-term solution, and quality control (QC) algorithms have been developed and implemented that succeeds in removing much of the RFI contamination from downstream products (Daniel Michelson, private communication).

However, aviation weather systems that use Canadian weather radar data to form continental scale mosaics, such as the Consolidated Storm Prediction for Aviation (CoSPA; Wolfson et al. 2008), NextGen Weather Processor (NWP; FAA 2022), and the planned Canadian Aviation Weather Systems (CAWS; NAV CANADA 2022), ingest a rawer form of data (“base data”) that lacks QC for mitigating RFI contamination. Consequently, RFI is routinely showing up in these systems. To address this issue, we developed a polarimetric QC procedure for censoring RFI contamination in the Canadian weather radar base data. This report describes its development, documents the detailed algorithmic steps, and shows some off-line test results. The RFI QC scheme has now been implemented in the NWP test reference system (TRS), and it is planned to be incorporated into CAWS as well.

This page intentionally left blank.

2. RFI CHARACTERISTICS

In order to maximize WD RFI detection probability and minimize the inadvertent censoring of weather data, we need to understand the characteristics of the offending RFI. To do this, we should examine all of the different observed variables available to us. CoSPA and NWP receive the following base data in Operational Program for the Exchange of Weather Radar Information (OPERA) Data Information Model (ODIM) Hierarchical Data Format, version 5 (HDF5) format (Michelson et al. 2019) from the Canadian weather radars.

- Reflectivity (horizontal polarization): Z_h (dBZ)
- Uncorrected reflectivity (horizontal polarization): Z_{uh} (dBZ)
- Radial velocity (horizontal polarization): V_h (m/s)
- Doppler spectrum width (horizontal polarization): W_h (m/s)
- Signal quality index (horizontal polarization): sqi_h (unitless)
- Differential reflectivity: Z_{dr} (dB)
- Co-polar correlation coefficient: ρ_{hv} (unitless)
- Differential phase: ϕ_{dp} (deg)
- Uncorrected differential phase: ϕ_{udp} (deg)
- Specific differential phase: K_{dp} (deg/km)

The uncorrected quantities are data that have not had clutter corrections applied to them (Leonardo 2021). Signal quality index (SQI) is the absolute value of the first lag in the signal autocorrelation divided by the power; it provides a measure of pulse-to-pulse signal coherence.

2.1 VISUAL EXAMINATION

First, let us examine the RFI characteristics visually through 2D range-azimuth plots. We take an example from the King City, Ontario (CASKR) site, collected on the 0.4° elevation scan at 00:05:15 UT on 12 November 2022. The polar coordinates are laid out in rectangular grid format so that near-range cells are equally resolved as the far-range cells. The data values are mapped to a color scale, and cells that were flagged as having missing values (i.e., where the product generation algorithm determined that the detection threshold for that variable was not met) are shown as blank (white).

Figure 2-1 shows Z_h . With CoSPA and NWP, this is the only variable that gets mosaicked to generate their output products. Therefore, this is the data field that must be cleaned of unwanted RFI contamination for reliable aviation weather products to be produced. This instance contains both extensive weather and RFI data. As expected, the WD RFI features look like vertical lines that are largely continuous if the signal is strong, but discontinuous (sometimes above and sometimes below the system’s detection threshold) when weak. In areas where RFI and weather signals coexist, the vertical-line nature of the RFI can become overshadowed by the underlying weather reflectivity patterns, depending on their relative signal strengths.

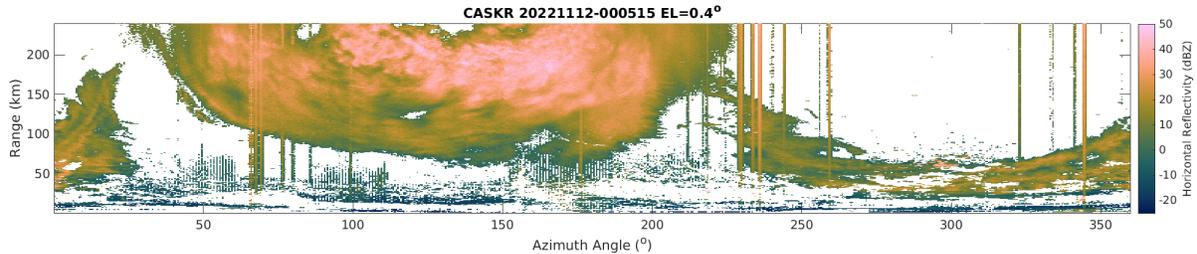


Figure 2-1. Range-azimuth reflectivity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4°.

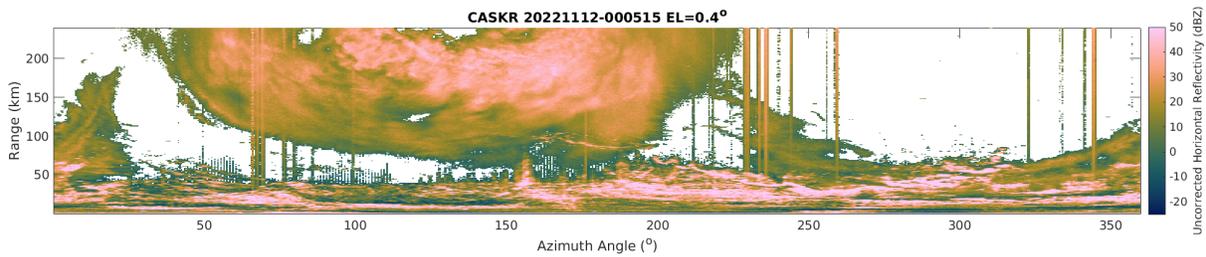


Figure 2-2. Range-azimuth uncorrected reflectivity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4°.

The uncorrected reflectivity (Figure 2-2) looks similar to Figure 2-1, except that the strong ground clutter at close range is not filtered out. Since this is not helpful in developing an RFI detection algorithm, we will drop consideration of Z_{uh} as an input from here on.

Next, we look at the other single-polarimetric base data. Figure 2-3 shows radial velocity. The RFI areas have noisy values, which is consistent with WD RFI manifesting as white noise in the Doppler spectral domain. It is also clear that the system detection threshold for velocity data is more stringent than for reflectivity, as more of the RFI data areas are missing compared to Figure 2-1. Similarly, the Doppler spectrum width plot (Figure 2-4) shows noisiness and data sparseness in the RFI areas. The sparsity of available data in areas where RFI is evident in reflectivity largely disqualifies V_h and W_h from consideration for RFI detection.

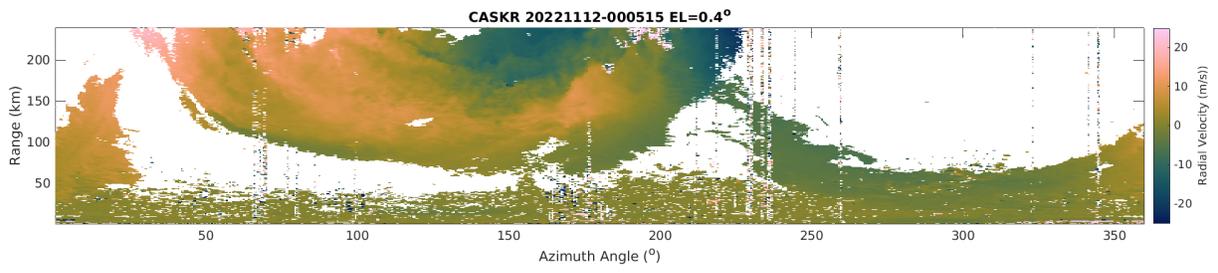


Figure 2-3. Range-azimuth radial velocity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4°.

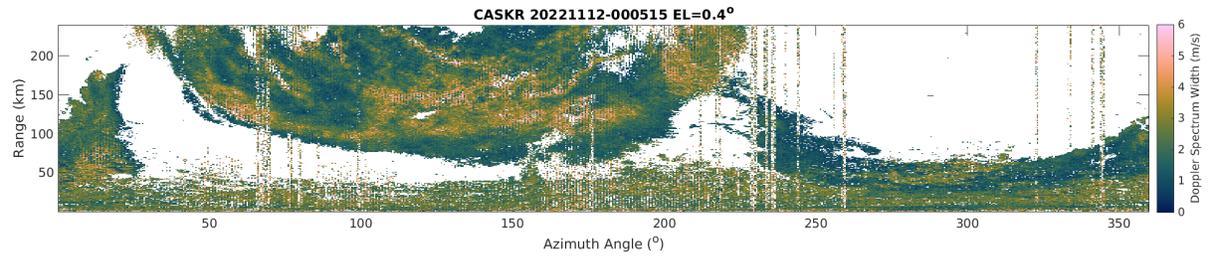


Figure 2-4. Range-azimuth Doppler spectrum width plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4°.

Moving on to the polarimetric base data, Figure 2-5 shows Z_{dr} . This variable exhibits a wide range of values in the RFI areas, with no obvious signature. Figure 2-6 shows ρ_{hv} . This variable appears intriguing for our purposes, since there is a clear tendency for RFI and ground clutter contaminated areas to have lower ρ_{hv} compared to weather-only areas.

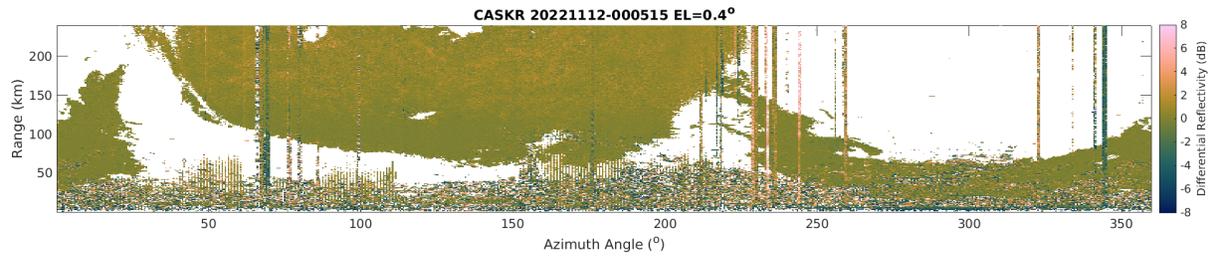


Figure 2-5. Range-azimuth differential reflectivity plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4°.

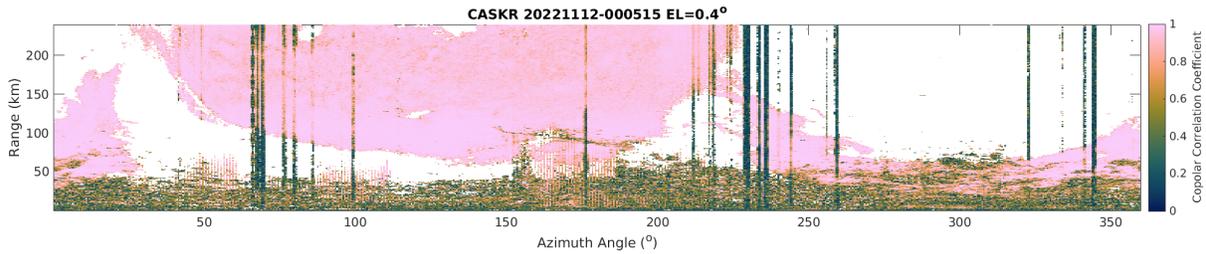


Figure 2-6. Range-azimuth co-polar correlation coefficient plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .

Figure 2-7 shows ϕ_{dp} . Apparently, the system detection threshold for differential phase is even more stringent than for radial velocity, as more of the RFI data areas are missing compared to Figure 2-3. On the one hand, we do not gain any information about the nature of RFI from this plot, but on the other hand, the editing “mask” for this variable may be a useful reference (or baseline standard) for RFI censoring. In other words, any RFI censoring scheme we develop must perform better than the differential phase detection threshold in both detection and false alarm metrics.

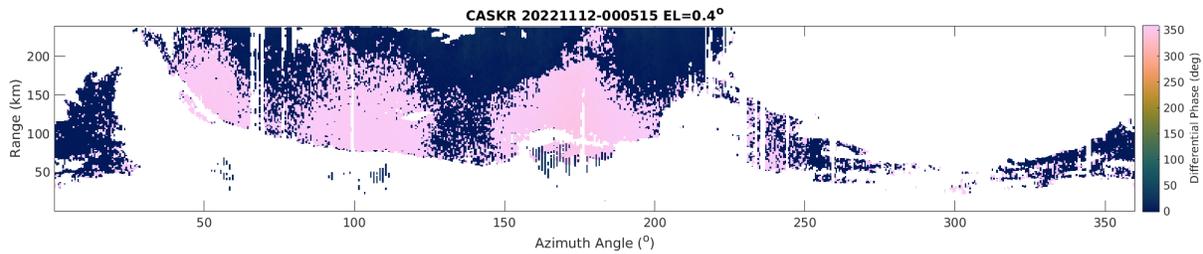


Figure 2-7. Range-azimuth differential phase plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .

The uncorrected version of differential phase is shown in Figure 2-8. The notable feature of this plot is that weather regions are smooth, while RFI and ground clutter contaminated areas are noisy.

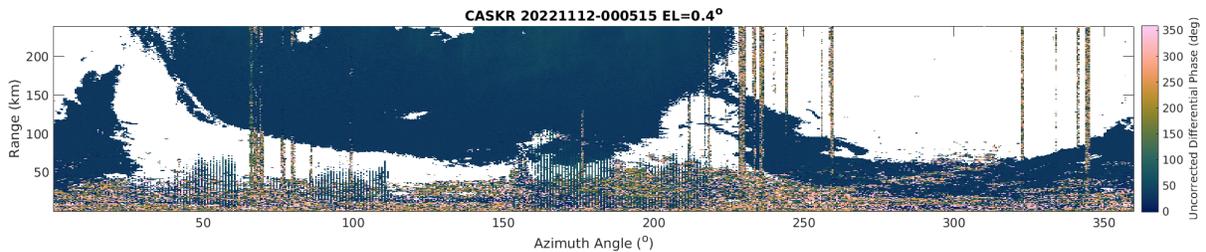


Figure 2-8. Range-azimuth uncorrected differential phase plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .

Figure 2-9 shows K_{dp} , which is generated by a least-squares slope fit on along-range ϕ_{dp} data (Leonardo 2021). As such, the editing mask (the white regions) is nearly identical to the one for ϕ_{dp} (Figure 2-7), and could also be used as a baseline for RFI censoring comparisons.

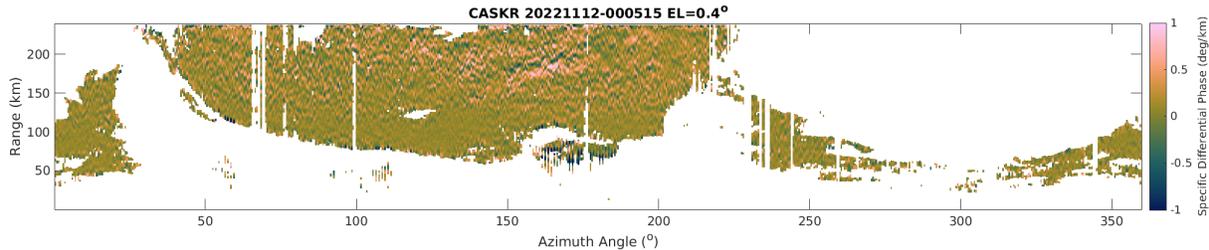


Figure 2-9. Range-azimuth specific differential phase plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .

Finally, we have sqi_h (Figure 2-10). This quantity is also of interest, because it is high for weather and low for RFI. (Its values for clutter residue appear to be noisy and wide-ranging.) It is also useful, because, unlike the other variables, it has not been thresholded, giving us full spatial coverage.

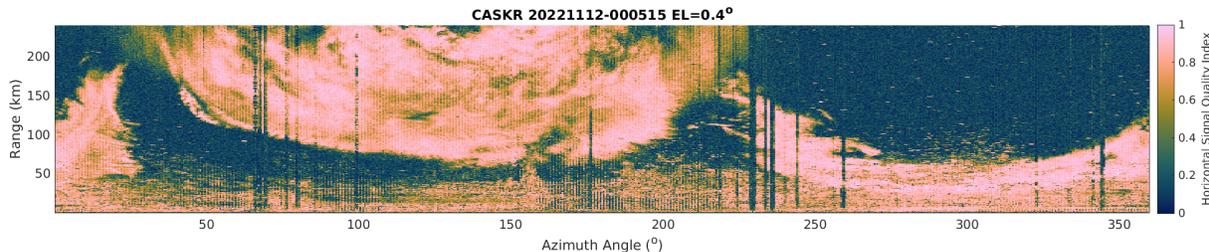


Figure 2-10. Range-azimuth signal quality index plot from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4° .

We also show the along-range variance of ρ_{hv} (Figure 2-11) and the along-range circular variance (Fisher 1993) of ϕ_{udp} (Figure 2-12), since they are potential candidates for discriminating between RFI and weather. In the plots, the variances are computed inside a moving window of ± 2 range gates. In both cases, the variance is high for RFI and clutter, and low for weather. Note that the along-range variance of ρ_{hv} indicates the presence of RFI within a weather patch more clearly than the along-range circular variance of ϕ_{udp} . This makes the former quantity a better candidate for initial identification of RFI-contaminated radials (first stage of algorithm as described in section 3.2) in the presence of weather.

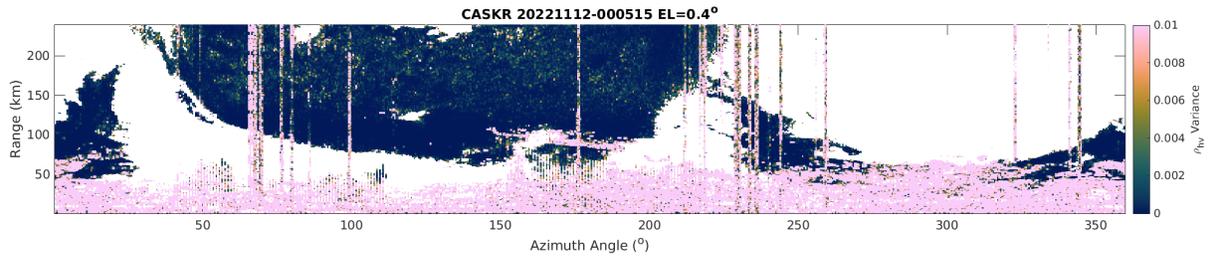


Figure 2-11. Range-azimuth plot of co-polar correlation coefficient along-range variance from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4°.

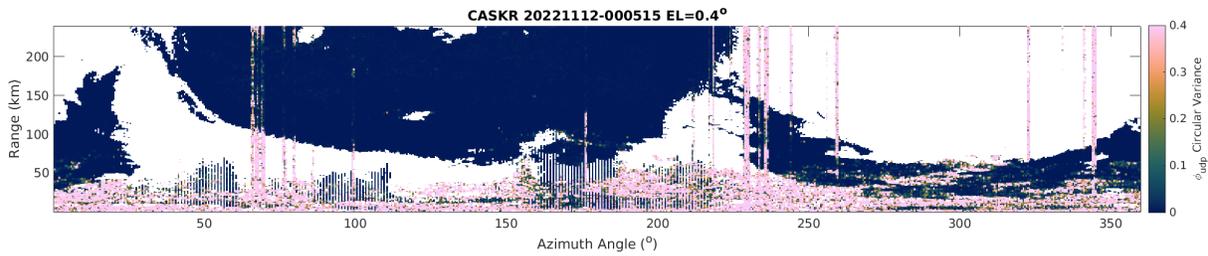


Figure 2-12. Range-azimuth plot of uncorrected differential phase along-range circular variance from the King City, Ontario weather radar at 00:05:15 UT on 12 November 2022, at an elevation angle of 0.4°.

At this point, we can narrow the focus to ρ_{hv} , sqi_h , the along-range variance of ρ_{hv} , and the along-range circular variance of ϕ_{udp} as the main candidates for WD RFI identification. We will now delve deeper by exploiting statistical analysis methods.

2.2 STATISTICAL ANALYSIS

In order to compare the statistical characteristics of weather vs. RFI data, we need clean (unmixed) cases for both types. For weather, we chose a scan from the Britt, Ontario (CASBI) site, collected on the 0.4° elevation scan at 22:05:15 UT on 25 January 2023 (Figure 2-13), and a scan from the Landrienne, Quebec (CASLA) site, collected on the 0.3° elevation scan at 00:41:05 UT on 12 November 2022 (Figure 2-14). As shown in these figures, only range gates greater or equal to 100 (50 km) were kept to exclude ground clutter contamination.

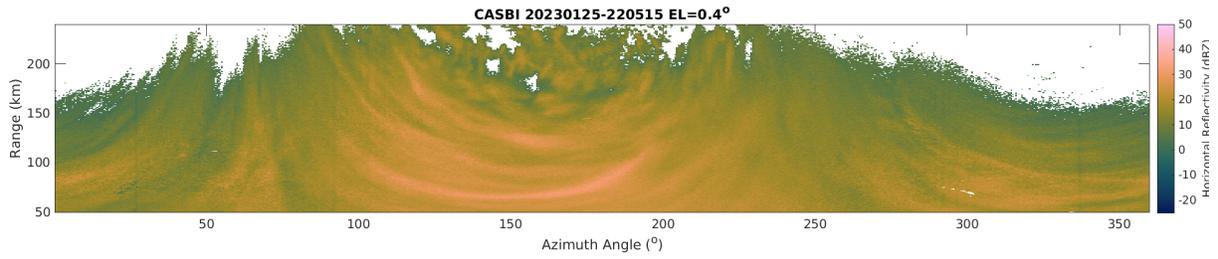


Figure 2-13. Range-azimuth plot of horizontal reflectivity from the Britt, Ontario weather radar at 22:05:15 UT on 25 January 2023, at an elevation angle of 0.4° .

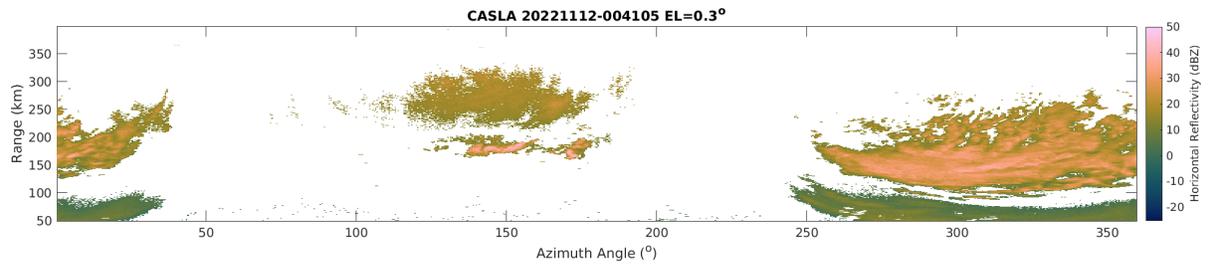


Figure 2-14. Range-azimuth plot of horizontal reflectivity from the Landrienne, Quebec weather radar at 00:41:05 UT on 12 November 2022, at an elevation angle of 0.3° .

For RFI, we chose two scans from the Franktown, Ontario (CASFT) site, collected on the 0.8° elevation scan at 14:58:42 UT (Figure 2-15) and the 0.4° elevation scan at 14:59:16 UT (Figure 2-16), both on 14 February 2023. Although there was no weather present, we only kept data from radials with RFI present as shown in the plots, in order to minimize contamination from clutter.

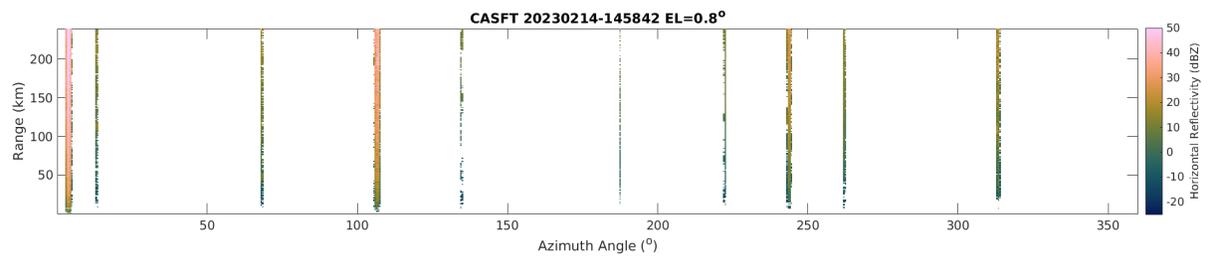


Figure 2-15. Range-azimuth plot of horizontal reflectivity from the Franktown, Ontario weather radar at 14:58:42 UT on 14 February 2023, at an elevation angle of 0.8° .

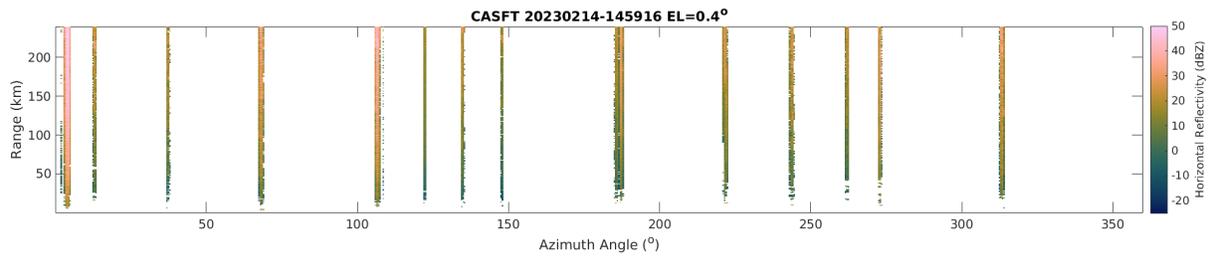


Figure 2-16. Range-azimuth plot of horizontal reflectivity from the Franktown, Ontario weather radar at 14:59:16 UT on 14 February 2023, at an elevation angle of 0.4°.

Figures 2-17 to 2-20 show histograms of co-polar correlation coefficient, SQI, co-polar correlation coefficient along-range variance, and uncorrected differential phase along-range circular variance for the weather and RFI case data. For all four quantities, the weather and RFI distributions are distinctly different, validating our visual examination of the range-azimuth plots.

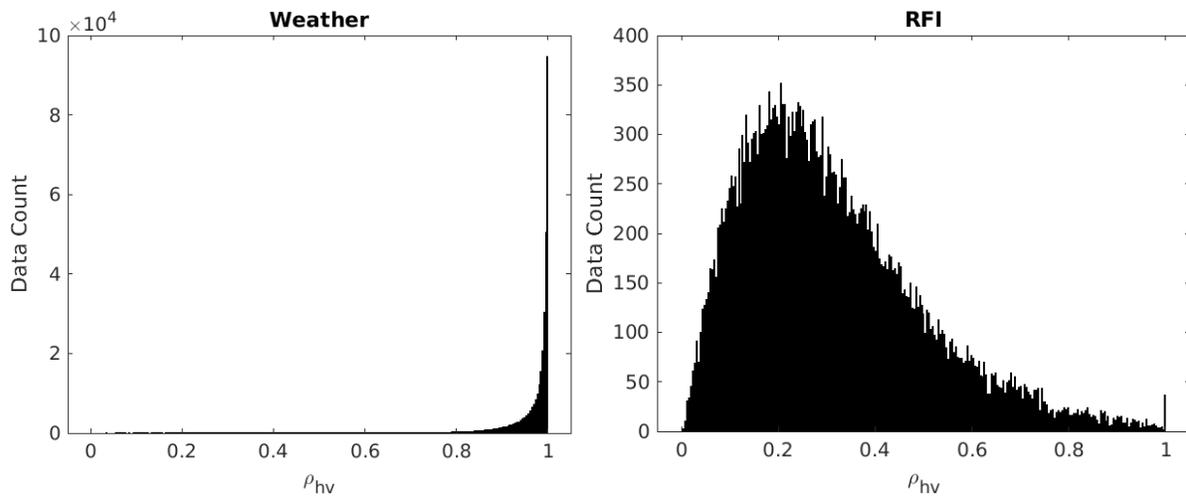


Figure 2-17. Histograms of co-polar correlation coefficient for the weather (left) and RFI (right) case data sets. The bin width is 0.02.

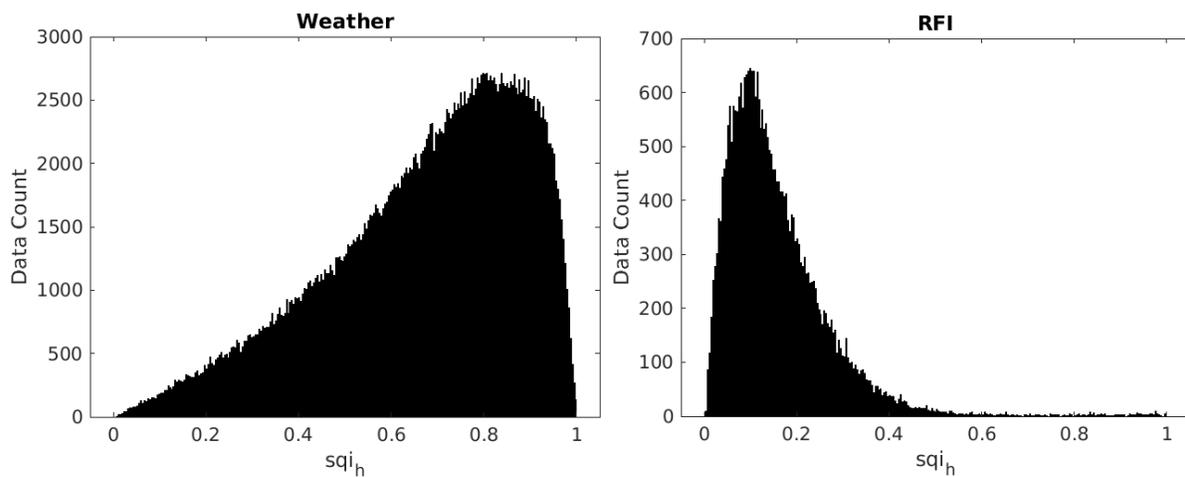


Figure 2-18. Histograms of SQI for the weather (left) and RFI (right) case data sets. The bin width is 0.02.

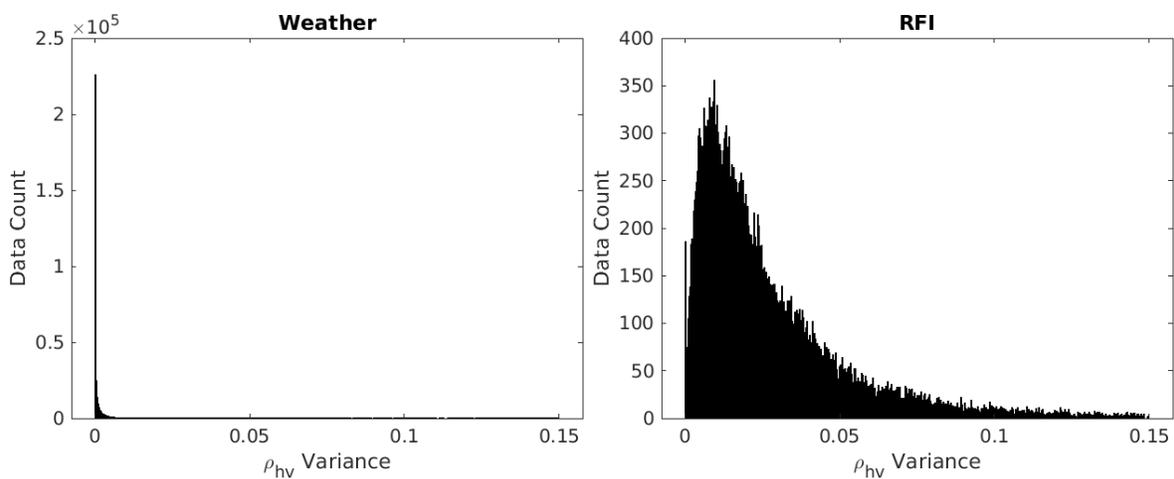


Figure 2-19. Histograms of co-polar correlation coefficient along-range variance for the weather (left) and RFI (right) case data sets. The bin width is 0.0003.

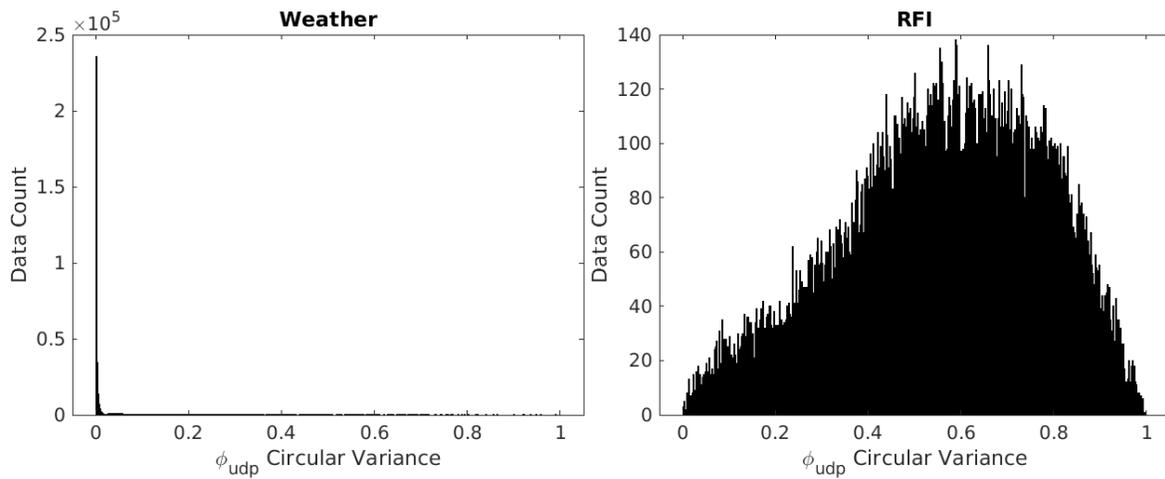


Figure 2-20. Histograms of uncorrected differential phase along-range circular variance for the weather (left) and RFI (right) case data sets. The bin width is 0.02.

From these histograms, we can compute receiver operating characteristic (ROC) curves for each of the four quantities. With co-polar correlation coefficient and SQI, the RFI distributions skew lower compared to the weather distributions. Thus, for these quantities, the percentage of data below a given threshold value yields the true positive rate for the RFI distribution and the false positive rate for the weather distribution. The opposite conditions hold for the variance quantities (i.e., percentage of data above a given threshold), because the RFI distributions skew higher relative to the weather distributions. The true and false positive rates covering the entire distribution span are then plotted to yield the ROC curves for these quantities (Figure 2-21).

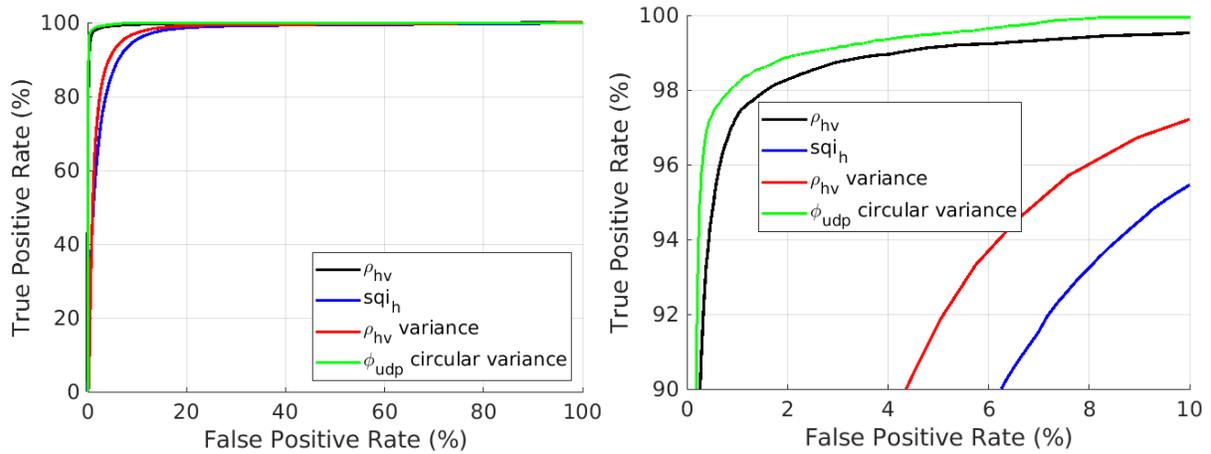


Figure 2-21. ROC curves for RFI detection against weather for co-polar correlation coefficient, SQI , co-polar correlation coefficient along-range variance, and uncorrected differential phase along-range circular variance. The left plot shows the full domain, while the right plot zooms in on the upper left-hand corner.

These ROC curves imply that ϕ_{udp} circular variance, followed closely by ρ_{hv} , should be the most effective discriminator between RFI and weather. Note, however, that the case data sets are not extensive, so these results should be taken with a grain of salt; there may be some distribution differences among various RFI sources and weather types. Nevertheless, these ROC curves are quite impressive, and they show great promise for an RFI discrimination algorithm to have excellent detection and false alarm performance. To be conservative, it may be advisable to utilize both ϕ_{udp} circular variance and ρ_{hv} in the algorithm.

The optimal discrimination threshold value (for the analyzed test data) can be determined from a ROC curve by finding the threshold value that minimizes the distance between the upper left-hand corner and the curves. The resulting RFI detection thresholds for these quantities are shown in Table 2-1. Again, these should be considered approximate values, since the input data sets were not extensive.

Table 2-1
Optimal RFI Detection Thresholds

Parameter	Criteria
ϕ_{udp} circular variance	> 0.084
ρ_{hv}	< 0.80
ρ_{hv} variance	> 0.0035
$sqih$	< 0.33

Finally, we note that the input data were carefully screened to exclude various non-RFI clutter contamination. Under normal operational conditions, even after ground clutter filtering, various residual clutter signals can be present, especially at low elevation angles where RFI signals tend to be strongest. Therefore, an RFI detector may sometimes flag these clutter cells as “RFI” (and other times not). In general, tagging clutter as RFI is not a bad thing, since clutter residue is undesirable. If a more thorough clean-up is needed, however, detectors specifically designed for each clutter residual type should be developed separately for maximum effectiveness.

3. RFI DETECTION AND CENSORING ALGORITHM

3.1 ALGORITHM FRAMEWORK

Previous work on RFI identification in weather radars (Cho and Frankel 2023) has shown that a direct range-azimuth cell-by-cell detection and censoring of RFI leads to too many false alarms (i.e., over-censoring of real weather data). Thus, as in the earlier work, we decided to adopt a two-stage approach in which radials that are contaminated with RFI are first identified, then, in the second stage, only those radials are checked gate-by-gate for RFI presence using the polarimetric variables.

Furthermore, after some initial algorithm development, we discovered that some RFI-contaminated radials were being missed occasionally. This can happen in cases when the RFI signal strength is such that (1) it is strong enough that the reflectivity data is not edited out, but (2) it is weak enough that the polarimetric data are edited out. (The radar’s signal processor employs different criteria for determining the “detection threshold” of each output variable.) To catch these instances, we decided to add a non-polarimetric RFI signature (“spike” along range) censor. Finally, to clean up stray non-weather cells, we added a speckle censor. A high-level flow diagram showing the entire process is given in Figure 3-1.

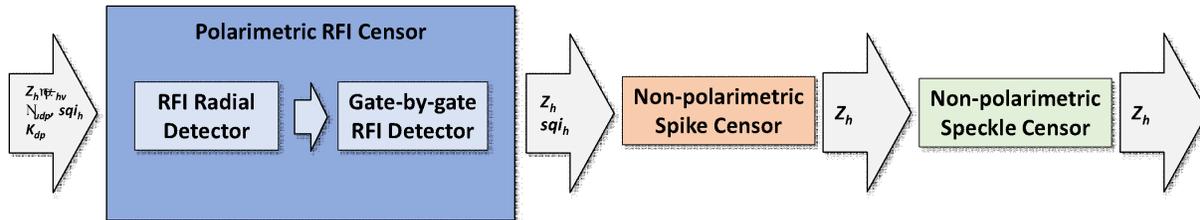


Figure 3-1. RFI detection and censoring flow diagram.

For our purposes (i.e., data quality control for CoSPA, NWP, and CAWS), RFI censoring is needed only on horizontal reflectivity data. For other applications, censoring may be required on the other base data output as well. Although this censoring algorithm may be used directly on the other base data as well, variable-specific tuning of the adjustable parameters might be required to optimize the trade-offs between detection and false alarm probabilities for each output data type.

3.2 POLARIMETRIC RFI CENSOR

For the first stage (determining RFI-contaminated radials), we experimented with various combinations of metrics and procedures. The one that worked the best overall was calculating the along-range variance of ρ_{hv} , multiplying it by the mean sqi_h over the variance-computation span, taking the median over all range gates, then thresholding. Multiplying by sqi_h decreases the contribution from ground clutter residue relative to RFI.

For the second stage (gate-by-gate RFI detection), we first check to see if K_{dp} has a valid value. If so, we move on to the next gate. Because, as noted in section 2.1, the radar’s signal processor “overcredits” K_{dp} such that essentially all RFI (plus some valid weather areas) are output as invalid, we can use this condition as an initial gateway to the rest of the polarimetric RFI censor in order to eliminate unnecessary computations. Then, to be conservative, we use the top two RFI/weather discrimination parameters (ρ_{hv} and along-range circular variance of ϕ_{udp}) in an “and” combination for RFI detection.

Here is a more detailed algorithm description:

The input variables are z_h , ρ_{hv} , sqi_h , K_{dp} , and ϕ_{udp} , all of array size $nGate \times nAz$. Note that the “below detection threshold” value of zero for ρ_{hv} and ϕ_{udp} is treated as a valid value in all relevant computations. The output consists of the z_h array edited for RFI.

The adjustable parameters for this algorithm and their current settings are:

- N_HALF_WINDOW_STAGE1 = 2
- RHOHV_VAR_MAX = 0.15
- SQI_DEF = 0.5
- RHOHV_RFI_THRES = 0.001
- N_HALF_WINDOW_STAGE2 = 2
- UPHIDP_VAR_THRES = 0.085
- RHOHV_MAX = 0.8

Stage 1: Check each radial for RFI contamination.

- a. Initialize logical 1D array $iRFIRadial$ of length nAz to false.

- b. In the input sqi_h matrix, replace any missing values with our default value of SQI_DEF. In the HDF5 source file, the value for missing sqi_h data is 255. (There are very rare instances of missing sqi_h data.)
- c. For each azimuth index, iAz , compute the local variance of ρ_{hv} , $rhohvVar$, at each range gate index, $iGate$:

$$rhohvVar(iGate, iAz) = \frac{1}{n-1} \sum_{iWin=iBeg}^{iEnd} [\rho_{hv}(iWin, iAz) - rhohvAve(iGate, iAz)]^2, \quad (3-1)$$

where

$$rhohvAve(iGate, iAz) = \frac{1}{n} \sum_{iWin=iBeg}^{iEnd} \rho_{hv}(iWin, iAz), \quad (3-2)$$

$iBeg$ = maximum of 1 or $iGate - N_HALF_WINDOW_STAGE1$,

$iEnd$ = minimum of $nGate$ or $iGate + N_HALF_WINDOW_STAGE1$, and

$n = iEnd - iBeg + 1$.

- d. For each azimuth index, iAz , compute the local mean of sqi_h , $sqiAve$, at each range gate index, $iGate$:

$$sqiAve(iGate, iAz) = \frac{1}{n} \sum_{iWin=iBeg}^{iEnd} sqi_h(iWin, iAz). \quad (3-3)$$

- e. Replace all values in the $rhohvVar$ matrix that are greater than RHOHV_VAR_MAX with 0. This is done to eliminate the effects of weather/non-weather boundaries on $rhohvVar$ that have nothing to do with RFI.
- f. Compute $rhohvRFI(iGate, iAz) = rhohvVar(iGate, iAz) \times [1 - sqiAve(iGate, iAz)]$ over all $iGate$ and iAz . Like RFI, ground clutter can have high $rhohvVar$, but its sqi_h tends to be higher than that for RFI, so using sqi_h (averaged for robustness) here helps differentiate between clutter and RFI.
- g. For each azimuth index, iAz , compute the median of $rhohvRFI$ over all $iGate$. (This gives a measure of how bad the RFI contamination—fraction of range gates and strength—is in the radial.) If the median value is greater than RHOHV_RFI_THRES, then set $iRFIRadial(iAz)$ to true.

Stage 2: Check each range gate for RFI in only the RFI-contaminated radials.

- a. For each azimuth index, iAz , proceed only if $iRFIRadial(iAz)$ is true.
- b. For each range gate index, $iGate$, proceed only if $K_{dp}(iGate, iAz)$ has an invalid value (-10), $z_h(iGate, iAz)$ has a valid value, and $rhohvAve(iGate, iAz)$ is less than $RHOHV_MAX$. The reflectivity value is used as a conditional here because an invalid value (-31.5) means that it was already edited out and so this cell does not need to be checked for the presence of RFI.
- c. Compute the local circular variance of ϕ_{udp} , $uphidpVar$:

$$uphidpVar(iGate, iAz) = 1 - \frac{1}{m} \sqrt{\left\{ \sum_{iWin=iBeg2}^{iEnd2} \cos[\phi_{udp}(iWin, iAz)] \right\}^2 + \left\{ \sum_{iWin=iBeg2}^{iEnd2} \sin[\phi_{udp}(iWin, iAz)] \right\}^2}, \quad (3-4)$$

where

$iBeg2 = \text{maximum of } 1 \text{ or } iGate - N_HALF_WINDOW_STAGE2,$

$iEnd2 = \text{minimum of } nGate \text{ or } iGate + N_HALF_WINDOW_STAGE2, \text{ and}$

$m = iEnd2 - iBeg2 + 1.$

- d. If $uphidpVar$ is greater than $UPHIDP_VAR_THRES$, then set $z_h(iGate, iAz)$ to the invalid flag value. The thresholds used for $rhohvAve$ and $uphidpVar$ were chosen to minimize false alarms (flagging cases where weather signal is dominant over RFI), while trying not to sacrifice detection probability. Further tweaking of the adjustable parameters is possible.

3.3 NON-POLARIMETRIC SPIKE CENSOR

Initially, we tried using the Next Generation Weather Radar (NEXRAD) Open Radar Product Generator (ORPG) spike editing algorithm (Smalley et al. 2008). This algorithm examines z_h data inside a moving 2D window of size $N_RANGE \times 2L + 3$ (range \times azimuth) to determine whether the N_RANGE gates in the central $2L + 1$ radials should be flagged for editing or not. This spike detection algorithm, however, had some drawbacks that needed to be addressed.

1. Practically all of the medium-to-strong RFI contamination instances are successfully detected by the polarimetric detector of section 3.2. However, some weak cases that fall below the detection thresholds for the polarimetric variables (yet not for reflectivity) are not flagged. (The polarimetric RFI detector may work even better if all the polarimetric variables were available

from the signal processor without editing.) These are the cases that need to be identified by the spike detector. With weak RFI, the reflectivity signal is not necessarily continuous in range, because the signal can fluctuate above and below the detection threshold for z_h . The problem is that the NEXRAD spike detection algorithm assumes that the spike reflectivity is unbroken in range. Thus, it is better in our case to parameterize the along-range z_h continuity using a fractional threshold to make it a less stringent requirement.

2. Sometimes valid weather features look like a spike to the detector, resulting in false removals. To help avoid these situations, we can use other information to decide if a spike candidate should be censored. The best candidate for our case is sqi_h , because, unlike ρ_{hv} and ϕ_{udp} , it is not edited below a detection threshold, i.e., in areas that it is most likely to be needed. Thus, apart from the rare missing data point, sqi_h is available for this purpose. If the algorithm were to be adapted for use on a different radar system that makes available ρ_{hv} and ϕ_{udp} data without editing, those would likely be better candidates given their superior performance in discriminating RFI from weather (Figure 2-21).
3. The NEXRAD spike algorithm does not scale well for $L > 1$. For example, the criterion that stipulates spike width from varying by more than L from gate to gate does not screen out cases where the spike is entirely to the left of center in one gate and entirely to the right in the next, since the spike width remains constant. While this may be acceptable for $L = 1$, it does not make sense for $L > 1$ —a spike should not be crooked. Also, the algorithm only detects and flags spikes of width $2L + 1$, so in order to catch spikes of different widths, it must be run multiple times with different L values. It would be more efficient to detect spikes of different widths in one pass.

In order to address these shortcomings, we developed a new spike detection and censoring algorithm. The algorithm description is as follows:

The input variables are z_h (edited by the polarimetric RFI censor) and sqi_h , both of array size $nGate \times nAz$. The output consists of the z_h array further edited for RFI.

The adjustable parameters for this algorithm and their recommended settings are:

- $L = 2$
- $N_RANGE = 10$
- $RANGE_FRAC_LIM = 0.35$
- $SQI_LIM = 0.3$

The new spike algorithm utilizes the same moving 2D window ($N_RANGE \times 2L + 3$) as the NEXRAD algorithm. The basic idea is to first determine whether each radial column of z_h within this window is “solid” or “sparse” (or neither). A radial is declared solid if the fraction of its range gates in the window with invalid z_h is less than BAD_WIN_FRAC . Conversely, a radial is declared sparse if the fraction of its range gates in the window with valid z_h is less than BAD_WIN_FRAC .

If the center radial is not solid, we do not set any flags and move the 2D window up by one range gate (or if the top of the window was the last range gate, move the 2D window to the next radial and back to the first range gate). If the center radial is solid, then we look for all contiguous radials that are also solid. If the contiguous solid radials are bounded on both sides by sparse radials, then all the gates inside the contiguous solid radials are considered a potential spike.

Finally, we check the potential spike region to see if it might have too much weather signal mixed in. To do this, we compute the mean sqi_h value over the spike region cells that have valid z_h values. If the mean sqi_h value there is less than SQI_LIM , then we set *spikeFlag* to true inside the spike region.

There are multiple ways to implement this algorithm. An implementation can certainly be different from the one outlined below as long as the outputs are the same.

1. Initialize *spikeFlag*, a logical 2D array of size $nGate \times nAz$, to false.
2. Define a moving 2D window of size $N_RANGE \times 2L + 3$ (range \times azimuth). Let the radial index, i , inside the window go from $-L - 1$ to $L + 1$, such that the middle radial corresponds to $i = 0$. This window is moved one range gate at a time, then one radial at a time, iteratively, until all range-azimuth cells in the radar scan are covered. (At the ends of the azimuth extents, wrap around to the other side as needed to fill in the window.) At each window position do the following:
 - a. Compute the fraction of 2D-window range gates in the center ($i = 0$) radial with invalid z_h (-31.5) values. If this fraction is less than $RANGE_FRAC_LIM$ (i.e., the center radial is solid), then proceed to step b. Otherwise, do nothing and move on to the next 2D window position.
 - b. Consider the left-hand ($i = -L - 1$ to -1) radials in the 2D window. Compute the fraction of range gates with valid z_h values for each of these radials. If at least one of the radials has a fraction less than $RANGE_FRAC_LIM$ (i.e., is sparse), proceed to step c. Otherwise, do nothing and move on to the next 2D window position.
 - c. Consider the right-hand ($i = 1$ to $L + 1$) radials in the 2D window. Compute the fraction of range gates with valid z_h values for each of these radials. If at least one of the radials has a fraction less than $RANGE_FRAC_LIM$ (i.e., is sparse), proceed to step d. Otherwise, do nothing and move on to the next 2D window position.

- d. Consider the sparse left-hand radials found in step b. Take the rightmost one (i.e., closest to the center radial), and check to see if all the radials between it and the center radial are solid (i.e., check to see if their fractions of range gates with invalid z_h values are less than RANGE_FRAC_LIM). If so, proceed to step e. Otherwise, do nothing and move on to the next 2D window position. (If the rightmost sparse radial is adjacent to the center radial, also proceed to step e.)
- e. Consider the sparse right-hand radials found in step c. Take the leftmost one (i.e., closest to the center radial), and check to see if all the radials between it and the center radial are solid (i.e., check to see if their fractions of range gates with invalid z_h values are less than RANGE_FRAC_LIM). If so, proceed to step f. Otherwise, do nothing and move on to the next 2D window position. (If the leftmost sparse radial is adjacent to the center radial, also proceed to step f.)
- f. Consider the radials in the 2D window lying between the rightmost left-hand sparse radial found in step d and the leftmost right-hand sparse radial found in step e. These solid radials form the candidate spike region. Compute the mean sqi_h value over the cells inside this spike region that have valid z_h values. If the mean sqi_h value is less than SQI_LIM, then set *spikeFlag* to true over the candidate spike region.
- g. After the previous steps have been followed for the entire radar scan data, set all z_h range-azimuth positions with *spikeFlag* set to true to the invalid z_h value.

The recommended setting of $L = 2$ allows the detection of RFI spikes up to five consecutive radials wide. If a site never experiences RFI that is more than three consecutive radials, then L can be reduced to 1. The trade-off between detection and false alarm probabilities can also be adjusted by tuning the value of RANGE_FRAC_LIM (higher for higher detection and false alarm probabilities, and lower for lower detection and false alarm probabilities).

Figure 3-2 provides illustrative examples of the spike detection logic. In the three 2D windows shown, range is vertical and azimuth is horizontal. Cells with valid z_h values are marked by “X.” Radials shaded yellow would have *spikeFlag* set to true if the mean sqi_h value computed over the yellow cells with “X” is less than SQI_LIM.

X		X	X		X	
		X	X	X		
X		X			X	
X			X			X
	X	X	X		X	
X		X	X			X
		X	X			
		X			X	
		X	X	X		
X		X	X			X

X			X		X	X
	X		X	X		
X		X	X		X	X
X			X			X
	X	X			X	
X			X			X
			X			
X	X		X		X	
	X		X			X
X	X		X		X	X

	X	X	X		X	
	X	X	X	X	X	
		X		X	X	
	X		X	X		
	X	X	X	X	X	
X		X		X		X
	X	X	X	X	X	
		X		X	X	
	X	X	X	X		
	X	X	X	X	X	X

Figure 3-2. Illustrative examples of spike detection using the recommended parameter settings.

3.4 NON-POLARIMETRIC SPECKLE CENSOR

A speckle remover applied at the end cleans up remaining isolated RFI cells (as well as isolated residue from various clutter sources). The algorithm used here is the NEXRAD ORDA speckle editor (Smalley et al. 2008). Briefly, a moving 2D window of size 5×5 (range \times azimuth) is centered on the cell of interest. If at least 75% of the neighboring cells within the window have invalid z_h values, then the center z_h value is set to invalid. This process is repeated three times over the entire radar scan.

This page intentionally left blank.

4. EXAMPLE RESULTS

First, we present the results of applying the three stages of the WD RFI censoring algorithm in succession, in order to show the contributions from each stage. Figure 4-1 is a fairly typical example in which the polarimetric RFI censor catches most of the offending areas, the spike censor cleans up the little bit of remaining RFI, and the speckle censor contributes negligibly. (Censored areas are shown in white.)

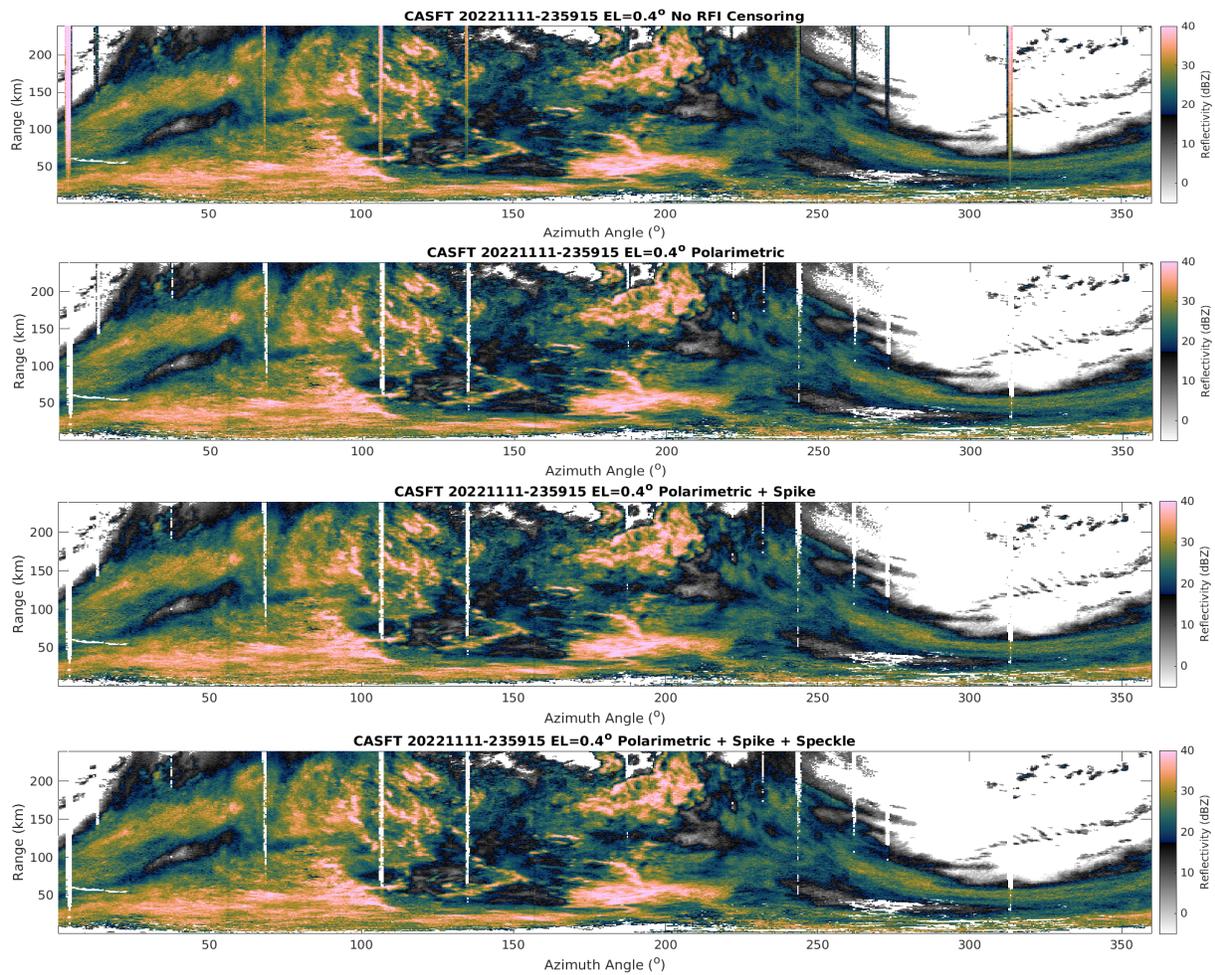


Figure 4-1. Range-azimuth plot of horizontal reflectivity from the Franktown, Ontario weather radar at 23:59:15 UT on 11 November 2022, at an elevation angle of 0.4°, without RFI censoring (top), with polarimetric RFI censoring (second from top), with polarimetric and spike censoring (second from bottom), and with polarimetric, spike, and speckle censoring (bottom).

Figure 4-2 shows one of the worst cases that we found for the polarimetric censor leaving a lot of RFI residue, which the spike and speckle censors effectively cleaned up in succession.

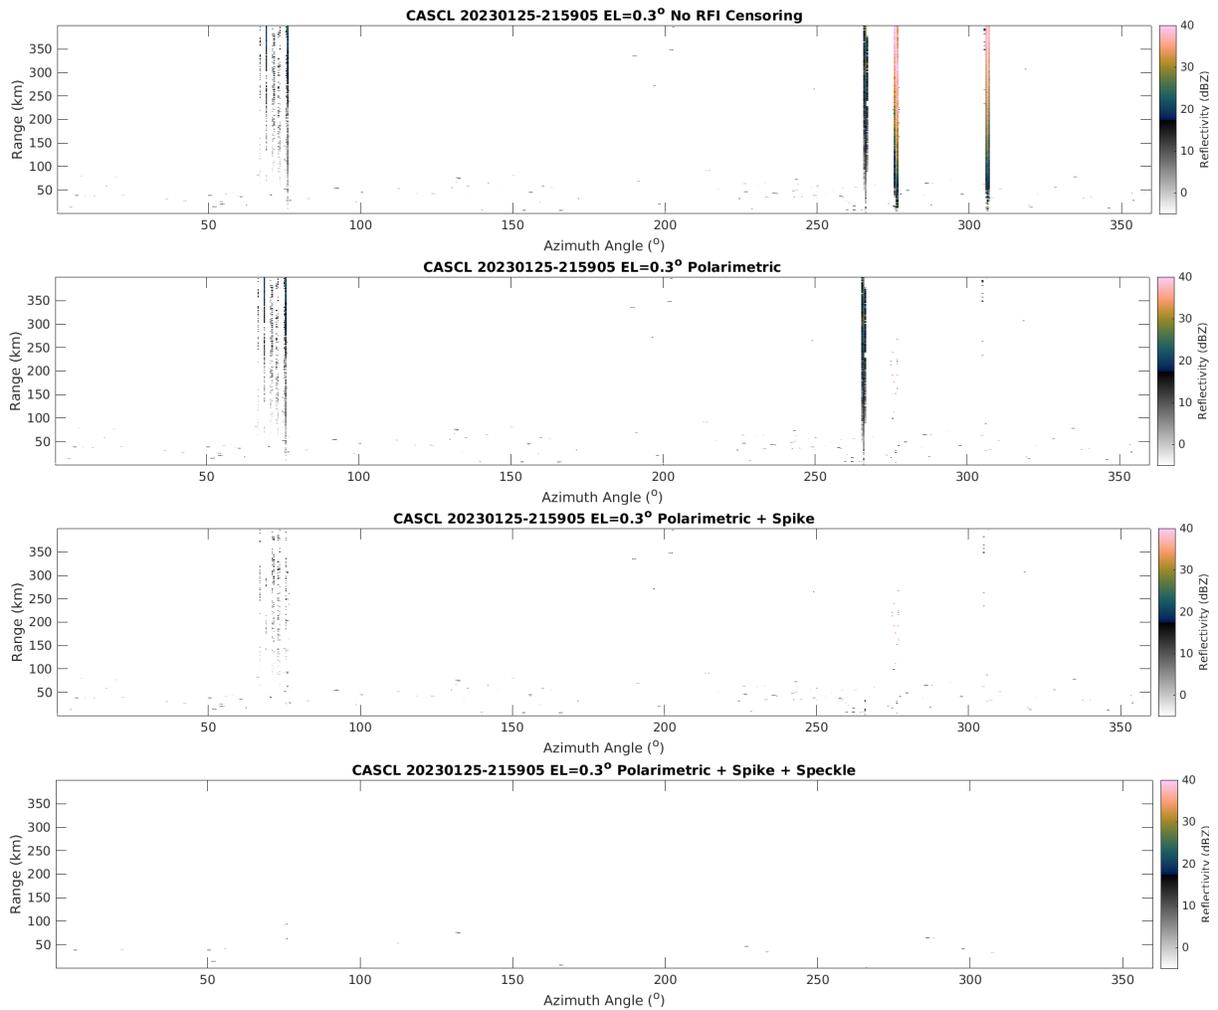


Figure 4-2. Range-azimuth plot of horizontal reflectivity from the Cold Lake, Alberta weather radar at 21:59:05 UT on 25 January 2023, at an elevation angle of 0.3°, without RFI censoring (top), with polarimetric RFI censoring (second from top), with polarimetric and spike censoring (second from bottom), and with polarimetric, spike, and speckle censoring (bottom).

Next, we compare the results of our RFI censoring scheme against simply using invalid K_{dp} data areas as an editing mask for reflectivity. Figure 4-3 shows a typical example when both RFI and weather are present. Clearly, the use of the K_{dp} mask would result in significant areas of valid weather being censored.

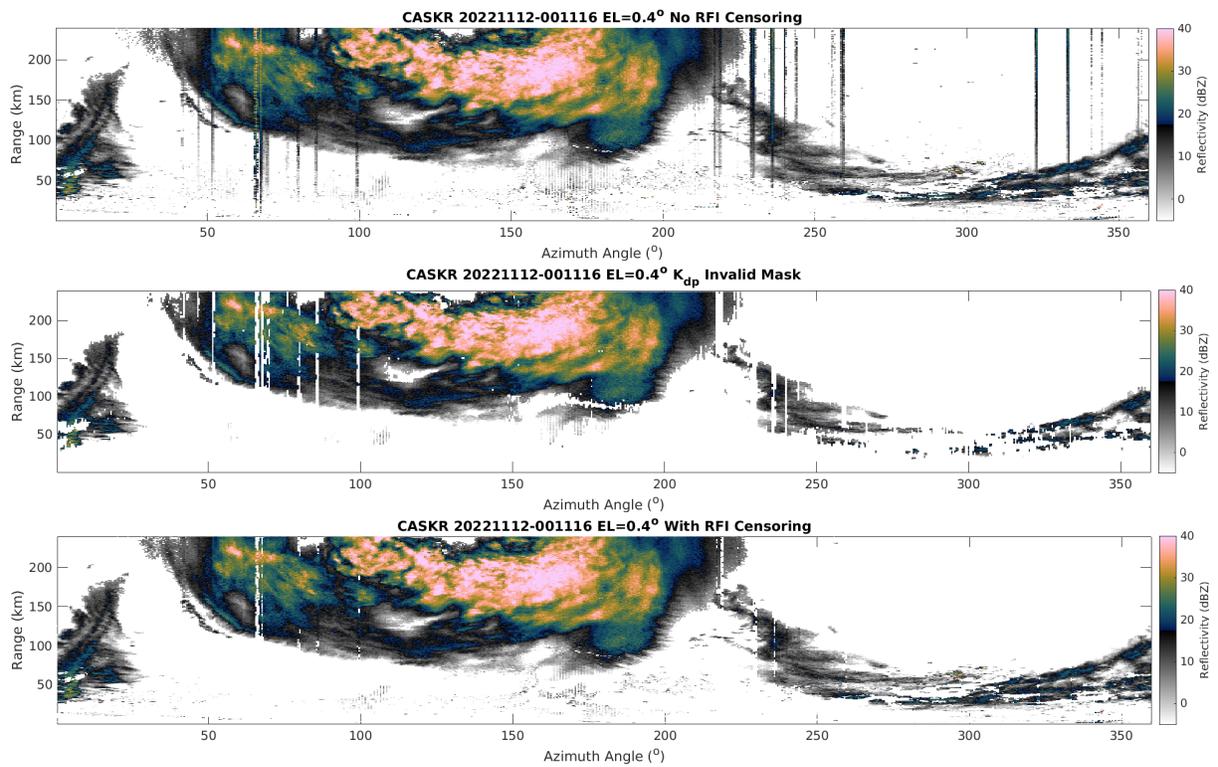


Figure 4-3. Range-azimuth plot of horizontal reflectivity from the King City, Ontario weather radar at 00:11:16 UT on 12 November 2022, at an elevation angle of 0.4° , without RFI censoring (top), with K_{dp} invalid data masking (middle), and with RFI censoring (bottom).

Figure 4-4 shows a case with only weather present. Again, it is clear that use of the K_{dp} mask would result in significant areas of valid weather being censored, while the RFI censoring algorithm preserves those areas.

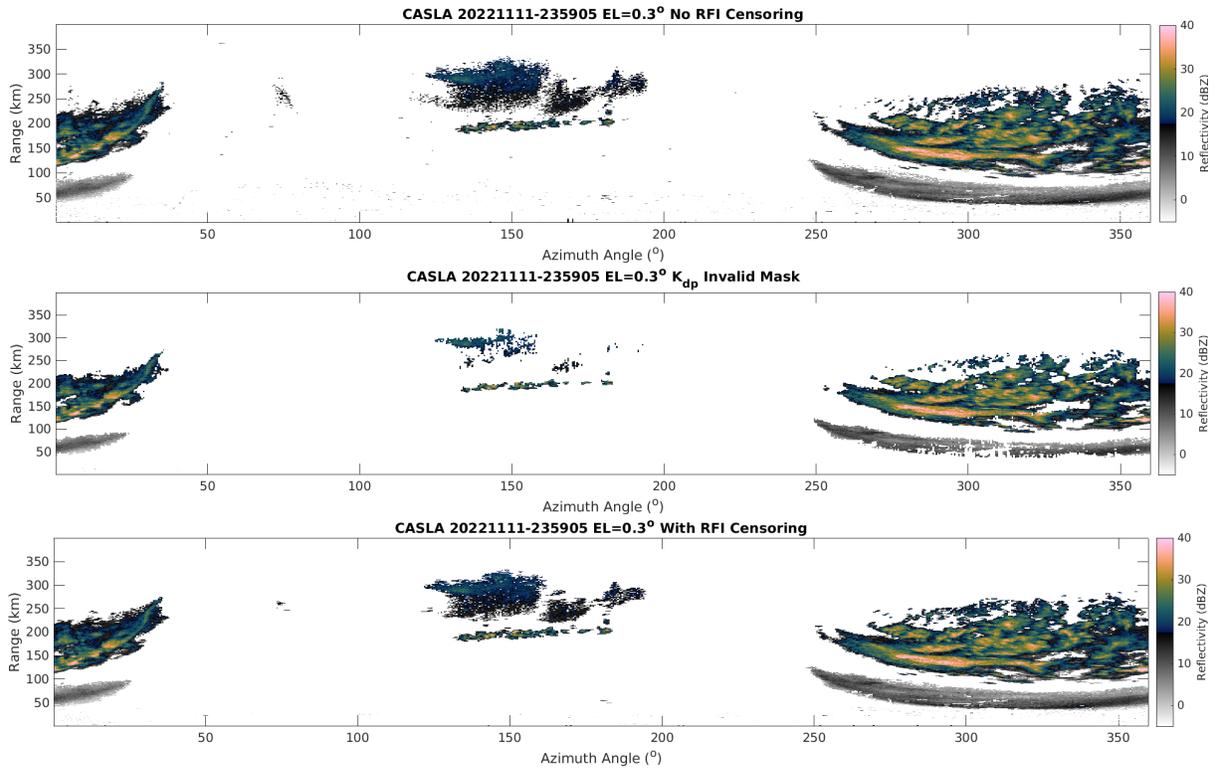


Figure 4-4. Range-azimuth plot of horizontal reflectivity from the Landrienne, Quebec weather radar at 23:59:05 UT on 11 November 2022, at an elevation angle of 0.3° , without RFI censoring (top), with K_{dp} invalid data masking (middle), and with RFI censoring (bottom).

Although these are limited examples, they are fairly typical cases, i.e., not cherry-picked to exaggerate algorithm performance. Extensive continuous testing on the real-time NWP TRS has shown that this new RFI censoring is effective and ready for operational implementation.

5. CONCLUSIONS

We developed an automated scheme to detect and censor WD RFI in reflectivity data output by the upgraded Canadian weather radars, based on the available set of base data. Extended real-time testing on the NWP TRS showed the performance of the RFI censoring scheme to be excellent. Based on the results, we recommend that the algorithm be implemented and deployed on relevant operational systems (NWP, CoSPA, CAWS) in the future.

Note that the development of the algorithm would have been more straightforward if the polarimetric base data from the radar's signal processor were not edited, i.e., the invalid data replaced with an invalid data flag value. A more computationally efficient (and potentially even more effective) algorithm might be possible if the polarimetric base data retained their original values everywhere with separate flag fields indicating the locations of invalid data. If the flag fields were encoded as single bits, the increase in data output rate may not be so egregious with respect to communication costs.

This page intentionally left blank.

GLOSSARY

1D	one-dimensional
2D	two-dimensional
CASBI	Britt, Ontario Canadian weather radar site
CASFT	Franktown, Ontario Canadian weather radar site
CASKR	King City, Ontario Canadian weather radar site
CASLA	Landrienne, Quebec Canadian weather radar site
CAWS	Canadian Aviation Weather Systems
CoSPA	Consolidated Storm Prediction for Aviation
HDF5	Hierarchical Data Format, version 5
NEXRAD	Next Generation Weather Radar
NWP	NextGen Weather Processor
ODIM	OPERA Data Information Model
OPERA	Operational Program for the Exchange of Weather Radar Information
ORPG	Open Radar Product Generator
QC	quality control
RFI	radio frequency interference
ROC	receiver operating characteristic
SQI	signal quality index
TRS	test reference system
UT	Universal Time
WD	wireless device

This page intentionally left blank.

REFERENCES

- Cho, J. Y. N., 2017: A new radio frequency interference filter for weather radars. *J. Atmos. Oceanic Technol.*, **34**, 1393–1406, <https://doi.org/10.1175/JTECH-D-17-0028.1>.
- Cho, J. Y. N., and R. S. Frankel, 2023: Radio frequency interference and anomalous velocity censoring algorithms for the Terminal Doppler Weather Radar. Project Memo. 43PM-Wx-0197, MIT Lincoln Laboratory, Lexington, MA, 27 pp.
- ECCC, 2023: Modernizing Canada’s weather-radar network. Environment and Climate Change Canada, Ottawa, Canada, <https://www.canada.ca/en/environment-climate-change/services/weather-general-tools-resources/radar-overview/modernizing-network.html>.
- FAA, 2022: NextGen Weather Processor. Federal Aviation Administration, Washington, DC, <https://www.faa.gov/nextgen/programs/weather/nwp>.
- Fisher, N. I., 1993: *Statistical Analysis of Circular Data*. Cambridge Univ. Press, New York, NY, 277 pp.
- Leonardo, 2021: GDRX signal processing user manual, document release 5.4. Document code GEMA-SD-0104-0116-0054, Leonardo Germany GmbH, Neuss-Rosellen, Germany.
- Michelson, D. B., R. Lewandowski, M. Szewczykowski, and H. Beekhuis, 2019: EUMETNET OPERA weather radar information model for implementation with the HDF5 file format, Version 2.3. European Meteorological Services Network, Brussels, Belgium.
- NAV CANADA, 2022: NAV CANADA partners with MIT Lincoln Laboratory on innovative weather decision support technologies. NAV CANADA, Ottawa, Canada, <https://www.navcanada.ca/en/news/blog/nav-canada-partners-with-mit-lincoln-laboratory-on-innovative-weather-decision-support-technologies.aspx>.
- Saltikoff, E., J. Y. N. Cho, P. Tristant, A. Huuskonen, L. Allmon, R. Cook, E. Becker, and P. Joe, 2016: The threat to weather radars by wireless technology. *Bull. Amer. Meteor. Soc.*, **97**, 1159–1167, <https://doi.org/10.1175/BAMS-D-15-00048.1>.
- Smalley, D. J., E. Mann, C. Ivaldi, and B. J. Bennett, 2008: Improving weather radar data quality for aviation weather needs. *13th Conf. on Aviation, Range, and Aerospace Meteorology*, New Orleans, LA, Amer. Meteor. Soc., 12.3, <https://ams.confex.com/ams/pdfpapers/132979.pdf>.
- Wolfson, M. M., W. J. Dupree, R. M. Rasmussen, M. Steiner, S. G. Benjamin, and S. S. Weygandt, 2008: Consolidated Storm Prediction for Aviation (CoSPA). *2008 Integrated Communications, Navigation and Surveillance Conf.*, Bethesda, MD, IEEE, 1–19, <https://doi.org/10.1109/ICNSURV.2008.4559190>.

This page intentionally left blank.

