

353 OPERATIONAL MITIGATION OF GROUND CLUTTER USING INFORMATION FROM PAST AND NEAR-FUTURE RADAR SCANS

Alexandra K. Anderson-Frey^{1*}, Frédéric Fabry²
¹Pennsylvania State University, University Park, PA
²McGill University, Montréal, QC

1. INTRODUCTION

When a radar pulse encounters obstacles in its path, the accuracy of radar reflectivity data is adversely affected, which in turn decreases the quality of forecasting and nowcasting tools such as rainfall totals and cell-tracking algorithms. The correction of these contaminated radar data presents a two-part problem to operational meteorologists and researchers seeking to apply these data quantitatively.

First, contaminated pixels must be identified as such: in the case of ground clutter (i.e., echoes produced by fixed objects close to the transmitter), this process involves identifying persistent clear-air echoes in the reflectivity imagery (Joss 1981) and using the fact that the power spectrum for ground clutter is localized around zero velocity to confirm identification using Doppler signal processing (Evans and Drury 1983). The second part of the problem consists of filling in these gaps caused by contaminated data, and it is this issue that will be addressed in this study.

Ground Clutter Mask for McGill Radar at 1.5 km

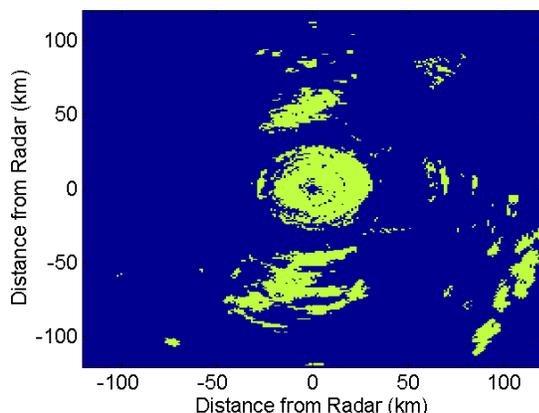


Figure 1: Depiction of the known stationary ground clutter at the McGill radar for the 1.5 km constant altitude plan position indicator (CAPPI) height. Locations of pixels contaminated by ground clutter are indicated in green, and each pixel represents a 1 km × 1 km horizontal square. The domain is 240 km × 240 km.

In order to evaluate the performance of the gap-filling algorithms to be developed in future sections, a number of test cases were studied at the S-band J.S. Marshall Radar Observatory in Sainte-Anne-de-Bellevue, some thirty km west of downtown Montreal,

**Corresponding Author Address:*

Alexandra Anderson-Frey, Department of Meteorology, The Pennsylvania State University, University Park, PA 16802; e-mail: aka145@psu.edu.

Quebec, Canada. Using the detection methods described above, a complete mask of ground clutter surrounding the McGill radar was obtained (Fig. 1).

The ultimate goal is to formulate an algorithm that will, with each new radar scan, automatically, quickly, and accurately replace each clutter-contaminated pixel with a value that more closely reflects the true meteorological situation. However, in order to evaluate the performance of each algorithm, it must be possible to compare the algorithms' output with some "true" reflectivity values. To this end, a realistic-looking region of "false" ground clutter will be created, where data has been blotted out in order to provide a testing ground for the gap-filling algorithms (Fig. 2).

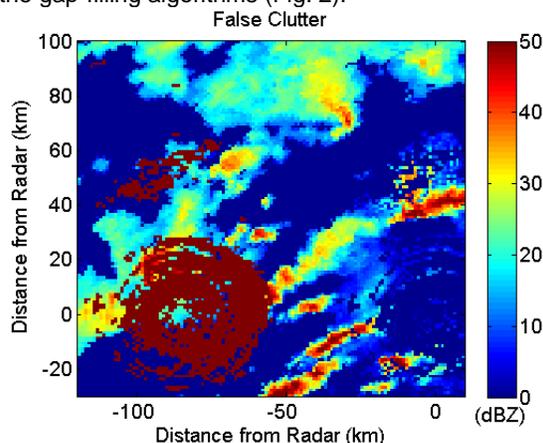


Figure 2: Example of the false ground clutter placed at the McGill radar for the 2.5 km CAPPI height, for use in testing the performance of gap-filling algorithms. The pixels in dark red represent areas that have been contaminated by ground clutter. (Note that the true ground clutter, in dark blue, is still present in the lower-right corner of the display, i.e., centered over the radar itself.)

The following section begins with the development of a gap-filling algorithm by addressing the simplest version of the problem at hand: given a single pixel contaminated by ground clutter, which pixel from its surroundings would make the best replacement (i.e., which pixel would produce the lowest error)? The algorithms in Section 3 expand this simplistic approach: rather than replacing a contaminated pixel with a single pixel from its surroundings, they use a combination of surrounding pixels, including pixels from radar scans at different heights and/or different times. Finally, Section 4 summarizes the results of the evaluation of these algorithms.

2. REPLACEMENT USING A SINGLE PIXEL

To begin, we consider a simplified version of the gap-filling problem: we have a single contaminated pixel that we seek to replace with a different single pixel from its surroundings, and we want to know what typical error we could expect to result from that replacement. A figure that depicts this "typical error" (in this case, the mean variance) is known as a variogram, and can be calculated by carrying out this replacement for every possible pixel in a given precipitation event, and then averaging over the number of replacements made. An example of a variogram, calculated using reflectivity data from the same convective event that provided the snapshot in Fig. 2, but excluding pixels declared as "contaminated" is shown in Fig. 3.

A variogram essentially provides a simple summary of the error structure of any given rainfall event. For example, in the event summarized by the variogram in Fig. 3, the fact that the variogram is stretched along the SW-NE axis indicates that, if given the choice between filling a gap using a pixel to the NE vs. a pixel to the NW, the reflectivity of the pixel to the NE will generally be closer to the true value of the contaminated pixel. This bias simply reflects this precipitation event's strong organization along the SW-NE diagonal: it makes more sense to replace a contaminated pixel with another pixel from the axis parallel to the squall line than with a pixel from the axis perpendicular to that squall line.

A simple gap-filling algorithm suggests itself immediately, based on these variograms: the "best pixel" approach will search through a list of available (i.e., non-cluttered) pixels and select the single pixel with the lowest variance to use as a replacement. Figure 4 demonstrates the value of using this best pixel approach with the same false ground clutter as that depicted in Fig. 2: the algorithm's output (left) compares quite well with the true reflectivity for that particular

radar scan (right). Even in the area denoted by the black ellipses, where the precipitation cells are almost entirely covered by the false ground clutter (Fig. 2), the qualitative appearance of the cells appears to have been captured by this best pixel approach.

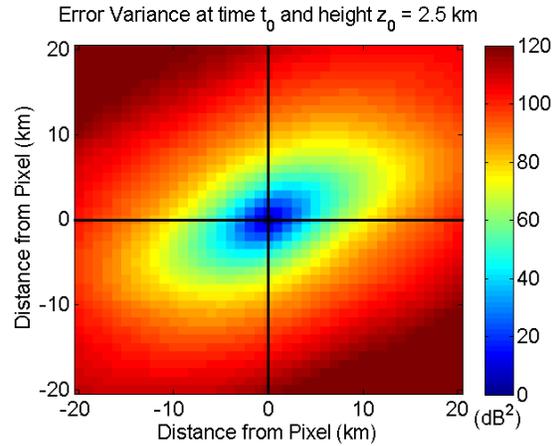


Figure 3: 2D variogram calculated for a convective precipitation event at a given time and height. This image can be thought of as a way to estimate how much error would result, on average, if we were to replace a reference pixel with a pixel from its surroundings. The crosshairs are centered on the location of the reference pixel. In terms of distance, 1 pixel = 1 km. Since the variance involves a sum of differences in reflectivities [dBZ], the units of this figure are dB^2 .

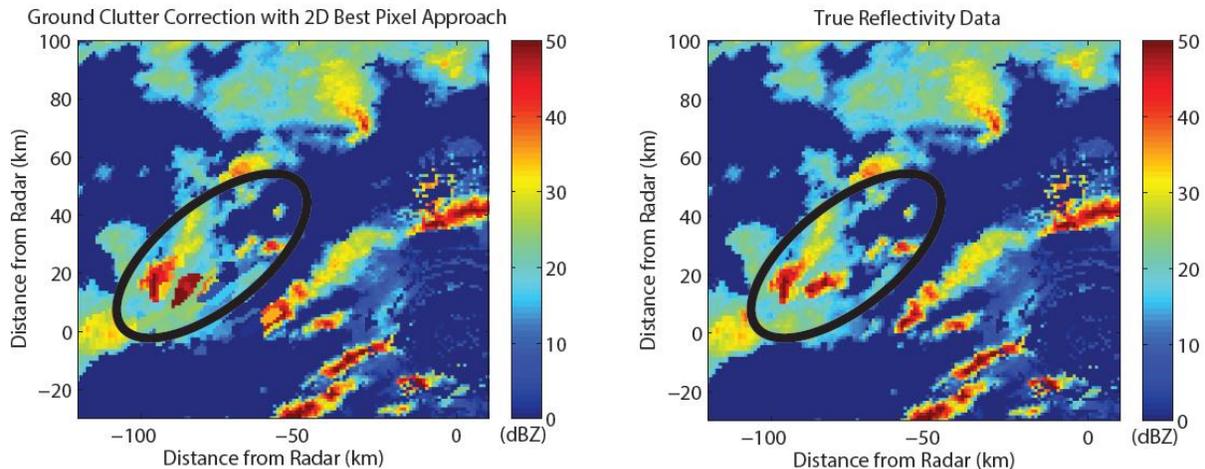


Figure 4: Comparison of the output produced by the two-dimensional best pixel approach (left) with the true reflectivity data (right) for a radar scan overlaid with the region of false ground clutter depicted in Fig. 2.

The limitations of this two-dimensional best pixel algorithm are clear: for one, the pool of possible replacement pixels is limited to those within the same radar scan as the contaminated pixel to be replaced. It seems likely, however, that a convective event with strong vertical development might see some value added by also considering pixels from different CAPPI levels, and it seems reasonable to surmise that data from different times (either earlier or, if product regeneration is feasible, in the near-future) would almost always provide at least some valuable information over the course of an entire precipitation event. In operational settings, we frequently ignore data from different radar scans, despite the fact that these data could contribute a large amount of valuable information.

In addition, the best pixel algorithm is focused on single-pixel replacement, whereas a combination of surrounding pixels could well produce a more accurate way to fill gaps left by ground clutter. The following section addresses this added complexity.

3. INFORMATION BLENDING

Rather than replacing each pixel contaminated by ground clutter with a single pixel from its surroundings, ground clutter can instead be replaced with some blend of data from its surroundings. One geostatistical method of blending information that lends itself particularly well to this problem is ordinary kriging (Krige 1951, Wackernagel 1995). A detailed overview of this method is outside the scope of this manuscript, but the general approach simply involves finding the mean of a number of pixels, as weighted by the error structure defined by the variograms (e.g., Fig. 3).

We are then faced with the problem of deciding how many pixels to combine (krige). Kriging a set of too many pixels will result in major computational delays (thus undermining the operational goals underlying this project), and will also tend to smooth out any extreme values, which is particularly undesirable given the importance of extreme values of reflectivity in severe weather forecasting. The number of pixels kriged in this study was chosen to be $2N$, where N is the number of dimensions of available potential replacement pixels, i.e., 2D limits the search for replacement pixels to a single radar scan, whereas 3D includes data from different heights, 3.5D also includes data from past times, and 4D also incorporates data from near-future times. This number of pixels to be kriged is low enough to reduce the risk of oversmoothing (and is well below the limit of computational power in an operational setting), but still represents a gain in available information for gap-filling algorithms.

The simplest technique that still makes use of ordinary kriging would thus be quite similar to the best pixel approach described in the previous section. Rather than simply picking the single pixel with the lowest variance, however, this approach would select the $2N$ pixels with the lowest variance, and would then combine them in a weighted average via ordinary kriging. The resulting value would then be used to replace the value of the pixel contaminated by ground clutter.

The dangers of this simple ordinary kriging method lie mainly in the fact that it allows for a great deal of redundant information: it will commonly select two adjacent, and hence very similar, pixels rather than picking just one of those pixels and then seeking out a different pixel that is more likely to provide independent information. In order to mitigate these pixel-selection issues, a "smart" ordinary kriging approach is developed to disallow the selection of adjacent or nearby pixels, all without increasing the number of pixels selected (and hence avoiding the computational cost and oversmoothing issues associated with ordinary kriging). By replacing redundant data with more independent information, this smart kriging process should fill the ground-clutter gaps with values that better reflect the mean pattern.

How should we define these "adjacent" or "nearby" pixels that are likely to contain redundant information, so as to avoid selecting them? Restricting ourselves to two dimensions for the sake of clarity (i.e., requiring that all four candidate replacement pixels are located on a single CAPPI scan from the same height and time), and picturing a set of orthogonal axes centered on the contaminated pixel that needs to be replaced, we can imagine selecting one pixel from along the positive x-axis, one from along the negative x-axis, one from along the positive y-axis, and one from along the negative y-axis. These four pixels are thus the combination that is least likely to sample redundant data, since they by definition sample four completely different quadrants of the radar scan.

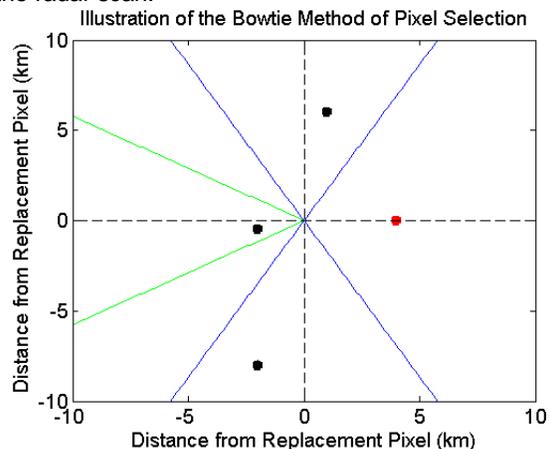


Figure 5: Illustration of the "bowtie" pixel selection method. The pixel with the lowest available variance (red point) is the first of the four pixels selected, and a system of axes (dashed lines) is built around it, based on the position of this best pixel with respect to the reference pixel to be replaced (origin). First, a 30-degree arc is measured out in the direction opposite that of the first pixel (green lines), and the pixel within that arc with the lowest variance is the next pixel to be selected (black point within the green lines). A similar methodology is used to select the remaining two pixels within the 30-degree arcs (blue lines) measured with respect to the orthogonal axis.

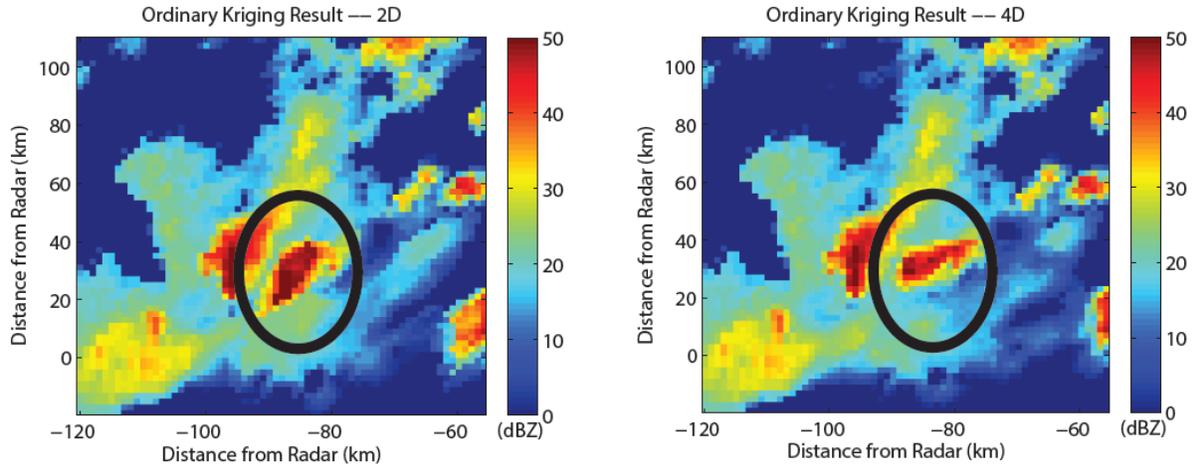


Figure 6: Qualitative depiction of the value added by including pixels from different heights and times. Nearly the entirety of the storm cell circled was blotted out by ground clutter (Fig. 2), and these figures show the results of the ordinary kriging algorithm with two different pools of potential replacement pixels: 2D only (left) and fully 4D (right).

In reality, however, this ambitious notion of picking pixels that align perfectly along each axis is unrealistic. We are not likely to be dealing with the replacement of a single contaminated pixel in isolation; large nearby regions of ground clutter dramatically narrow the pool of candidate pixels to choose from. By widening the candidate pixel search parameters from "must be located on the axis" to "must be located within 30 degrees of the axis," we can still find four uncluttered replacement pixels, and simultaneously ensure that they are likely to provide relatively independent information. This strategy lays the foundation for the "bowtie" method of pixel selection (Fig. 5).

This process is generalizable to four dimensions (using statistical equivalents between variances in time and horizontal distance), and hence can take advantage of the massive store of valuable information provided by reflectivity data at different CAPPI heights and times. Figure 6 illustrates the qualitative value added, first by switching from the best pixel approach to the ordinary kriging approach (left; compare with Fig. 4, noting that the zoom level is different), and then by switching from a simple 2D approach to a fully 4D approach (right). When compared with the true intensity and size of this cell (Fig. 4), the 4D algorithm is a very close match.

4. RESULTS AND CONCLUSIONS

Table 1 summarizes the results for each of the four gap-filling algorithms. The nearest neighbor approach simply replaces each contaminated pixel with its nearest uncontaminated neighbor, which is approximately the level of complexity examined in other work studying operational gap-filling in contaminated radar data (e.g., Lee et al. 1995, Bellon and Kilambi 1999, Galli 1984, Sánchez-Diezma et al. 2001). The best pixel approach is as described in Section 2, and the two ordinary kriging approaches are described in Section 3. Note that there are two vectors of improvement: more accurate

gap-filling can be achieved either by increasing the complexity of the pixel-replacement algorithm (i.e., going from a best pixel approach to a simple ordinary kriging approach to a smart ordinary kriging approach), or by increasing the dimensionality, and hence the pool of possible replacement pixels.

All of the algorithms introduced in this work outperform the simple nearest-neighbor approach, but the greatest leap in performance occurs across the board when progressing from a three-dimensional algorithm to one that includes data from other times (even if only past times). In particular, the incorporation of data from different times improves performance most dramatically for the smart ordinary kriging approach that uses the "bowtie" pixel selection method (Fig. 5).

While this work has primarily been a proof-of-concept for the operational use of the smart ordinary kriging algorithm in ground clutter mitigation, its implications extend to other areas of radar meteorology. First, the ordinary kriging method itself is a simple and versatile technique that has seen very little use in radar meteorology, having mainly been limited to interpolation between broadly spaced data points, such as rain gauge networks (Atkinson and Lloyd 1998). This is not particularly surprising, since denser datasets such as radar require far more computational power for ordinary kriging: essentially, users of radar data have too much information to effectively make use of this powerful method. Through the use of the bowtie pixel-selection method described in the previous section, however, this problem is rendered moot: by selecting a small number of representative pixels that can then be kriged for relatively little computational cost, the range of feasible applications for this method is broadened.

Table 1: Mean standard deviations of radar reflectivity (dB; essentially, mean error) for each of the gap-filling algorithms used in the convective precipitation event illustrated in the preceding figures.

Nearest Neighbor	
2D	5.3
Best Pixel	
2D	5.3
3D	5.2
3.5D	4.9
4D	4.9
Ordinary Kriging	
2D	4.4
3D	4.4
3.5D	4.1
4D	4.0
Smart Ordinary Kriging	
2D	4.3
3D	4.3
3.5D	3.0
4D	2.9

Second, the applications of these radar data gap-filling algorithms need not be limited to ground clutter. While ground clutter was the focus of this study, being particularly easy to identify and map out ahead of time, there are also complex algorithms in place to identify, say, regions of radar attenuation in real-time (Gorgucci et al. 1998). There is no reason why, once attenuated pixels have been identified, the user's choice of ordinary kriging algorithm could not then be used to immediately fill in these newly identified gaps. Gap-filling in other fields, so long as those fields behave at least somewhat similarly to radar reflectivity (i.e., they are continuous), could also benefit from this approach. This smart ordinary kriging algorithm could also be used to, for instance, fill gaps in Doppler velocity imagery.

Finally, and perhaps most importantly, in a field that frequently treats individual radar scans as though they have just sprung into being in perfect isolation from all that has preceded them, this study has provided a proof-of-concept for the importance and value of considering data from different times. Data from earlier radar scans is a relatively untapped well of information, and given the substantial improvements in even these simple gap-filling approaches, it is well past time that radar products should be generated and applied with full knowledge and appreciation of the data that came before.

6. REFERENCES

- Atkinson, P.M., and C.D. Lloyd, 1998: Mapping precipitation in Switzerland with ordinary and indicator kriging. *J. Geog. Info. Dec. Anal.*, **2**, 65-76.
- Bellon, A., and A. Kilambi, 1999: Updates to the McGill RAPID (Radar Data Analysis, Processing and Interactive Display) system. Preprints, *29th Int. Conf. on Radar Meteorology*, Montreal, Amer. Meteor. Soc., 121-124.
- Evans, J.E., and W.H. Drury, 1983: Ground clutter cancellation in the context of NEXRAD. Preprints, *21st Conf. on Radar Meteorology*, Edmonton, Amer. Meteor. Soc., 158-162.
- Galli, G., 1984: Generation of the Swiss radar composite. Preprints, *22nd Conf. on Radar Meteorology*, Zurich, Amer. Meteor. Soc., 208-211.
- Gorgucci, E., G. Scharchilli, V. Chandrasekar, P.F. Meischner, and M. Hagen, 1998: Intercomparison of techniques to correct for attenuation of C-band weather radar signals. *J. Appl. Meteor.*, **37**, 845-853.
- Joss, J., 1981: Digital radar information in the Swiss Meteorological Institute. Preprints, *20th Conf. on Radar Meteorology*, Boston, Amer. Meteor. Soc., 194-199.
- Krige, D.G., 1951: *A statistical approach to some mine valuations and allied problems at the Witwatersrand* (MS thesis). University of Witwatersrand, Johannesburg, South Africa.
- Lee, R., G. Della Bruna, and J. Joss, 1995: Intensity of ground clutter and of echoes of anomalous propagation and its elimination. Preprints, *25th Conf. on Radar Meteorology*, Paris, Amer. Meteor. Soc., 651-652.
- Sánchez-Diezma, R., D. Sempere-Torres, G. Delrieu, and I. Zawadzki, 2001: An improved methodology for ground clutter substitution based on a pre-classification of precipitation types. Preprints, *30th Int. Conf. on Radar Meteorology*, Munich, Amer. Meteor. Soc., 271-273.
- Wackernagel, H., 1995: *Multivariate Statistics—An Introduction with Applications*. Springer, Berlin, 389 pp.