Operational mitigation of ground clutter using information from past and near-future radar scans

Alexandra Anderson-Frey

Department of Atmospheric and Oceanic Sciences

McGill University, Montreal

August 2013

A thesis submitted to McGill University in partial fulfilment of the requirements for the degree of Master of Science

Abstract

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. In this study, we seek an optimal solution for real-time, operational gap-filling in radar data contaminated by known areas of ground clutter, and explore a variety of algorithms of increasing complexity to that end, making use of a geostatistical method known as ordinary kriging. The final result is the development of a "smart" ordinary kriging algorithm. This method replaces clutter-contaminated pixels in radar data using the weighted average of a nearby collection of uncontaminated pixels, which have been specially selected to sample independent spatial and temporal information while avoiding bogging down calculations with redundant information. These data are obtained not only from the same reflectivity scan as the ground clutter to be corrected, but also from different heights and from both earlier and near-future times. The incorporation of the time dimension in particular adds a great deal of value to simplistic algorithms, even when only data from past times are considered. Radar scans from earlier times are thus shown to be a major untapped source of information that can be used to generate (and regenerate, using near-future data) more accurate radar products.

Résumé

Lorsqu'un signal radar rencontre des obstacles, la précision des données de réflectivité est endommagée, ce qui réduit la qualité des outils de prévisions météorologiques tels les totaux de précipitation et les algorithmes qui surveillent l'évolution des orages. Dans cette étude, on recherche une solution optimale pour remplir en temps réel et dans un contexte opérationnel les trous d'information causés par les échos de sol, en explorant une variété d'algorithmes de plus en plus complexe basée sur une méthode géostatistique: le kriging ordinaire. Le résultat final est le développement d'un algorithme de kriging ordinaire "intélligent". Cette méthode remplace les pixels contaminés en utilisant la moyenne pondérée de pixels non-contaminés à proximité, où ces pixels sont sélectionnés specialement pour incorporer des données indépendentes et pour ne pas surcompliquer les calculs avec trop d'informations redondantes. Ces informations proviennent non seulement du même temps et du même niveau que la région qui doit être corrigée, mais aussi d'aux autres niveaux ainsi que du passé et du proche-futur. L'inclusion de la dimension temporelle en particulier offre grand valeur même aux algorithmes les plus simples, et aussi lorsqu'on considère seulement les informations du passé et non du futur. Les données du radar des temps antérieurs constituent alors une source inexploité d'informations

qui pourraient permettre de générer (et de régénérer, en utilisant les données du proche-futur) des produits radar plus précis.

Acknowledgements

I would like to acknowledge and thank my advisor, Prof. Frederic Fabry, for his unflagging support and boundless patience throughout the sometimes rocky preparation of this thesis. His advice, suggestions, encouragement, and stable of tips and tricks for scientific writing and presentation have been instrumental every step of the way. I would also like to thank the research group at the J.S. Marshall Radar Observatory, especially Aldo Bellon, for his data, suggestions, and support. The radar group at McGill have been excellent sounding boards for research ideas and theories, and many discussions with Madalina Surcel and Ya-Chien Feng sparked ideas for new programming or research techniques.

Finally, I would like to acknowledge and thank my family for their love and understanding from day one, no matter how baffling my passion for this research undoubtedly seemed to them at the time.

Dedication

To everyone I lost along the way: to Gracie and Midnight for their constant and frequently ridiculous companionship, to Bob and Edie Anderson and Bert Frey for their confidence and love, and to Mike Denton, who never saw science as anything but an adventure.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

1.1 Operational radar meteorology and data contamination

Any forecaster, researcher, or member of the public who has spent any amount of time watching radar output will likely have noticed some data that do not represent the true meteorological state of the atmosphere (e.g., data that are affected by ground clutter, attenuation, beam blockage, etc.). If, for instance, a small but powerful thunderstorm is completely obscured by these bad pixels, the forecast of the storm's track and intensity will be contaminated—as will products such as one-hour accumulation totals, which in turn will pass this error along to hydrological and climatological applications. This data contamination issue is complex, but attempts to resolve it are worthwhile given its implications in a variety of important applications.

The correction of 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, which can be an extremely difficult proposition in itself, given the rapid evolution and complexity of meteorological fields. Second, once identified, these contaminated pixels must be either removed from consideration or replaced with data that better reflect the true meteorological situation. These two topics have received a wealth of attention in the literature, and a brief summary of each part of the problem will be presented in the following sections. After this flyover, the remainder of the study will be devoted to the optimization of an algorithm designed to tackle the second part of the problem (i.e., correction of radar data given a known array of contaminated pixels) in real time.

1.2 Identification of contaminated pixels

In order to correct bad pixels, it is first necessary to define what constitutes a "contaminated" pixel in the first place. The identification of three types of contaminated pixels will be briefly explored in this section in order to provide context for the development of the ground clutter correction algorithm that follows.

One area that has seen a great deal of interest is that of attenuation, the process whereby the energy of a radar pulse is reduced during its passage through regions of intense precipitation (or even a film of water on the radome itself), resulting in anomalously low reflectivity values further from the radar. Figure 1.1 shows a dramatic example of such attenuation as measured by the OU-PRIME radar (Palmer et al. 2011). The reflectivity imagery (Fig. 1.1a) strongly hints at the existence of a region of attenuation beyond the half-circle of high reflectivity associated with large hydrometeors advected around a mesocyclone. The differential propagation phase Φ_{DP} (Fig. 1.1b) shows this attenuation still more clearly as streaks originating in the same location as the half-circle of higher reflectivity. All data obscured by this region of higher reflectivity are thus suspect, and some level of correction is necessary before using these data in quantitative studies.

Not all regions of attenuation are as obvious as those shown in Fig. 1.1, however. Gorgucci et al. (1998) caution that at C-band wavelengths, typical attenuation is often small enough that it is not easily observable, and yet it is large enough to affect quantitative analysis of reflectivity data. In addition, Scarchilli et al. (1993) note that attempting to correct relatively small C-band attenuation (~1 dB) introduces error greater than the effects of the attenuation itself. These and similar issues have forced researchers to adopt a variety of creative approaches to the problem of C-band attenuation identification and correction.

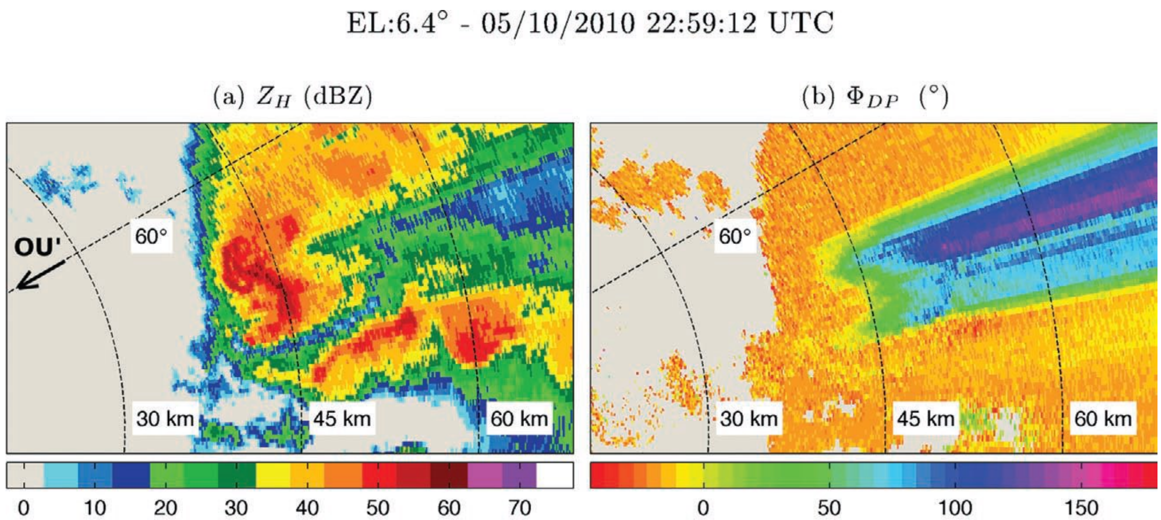

Figure 1.1: Examples of radar products contaminated by attenuation. The range rings indicate distance in km from the OU-PRIME radar. (a) Radar reflectivity at a 6.4° elevation scan shows a half-circle of > 50 dBZ reflectivity occurring at a distance of approximately 45 km from the radar, with much weaker (< 10 dBZ) reflectivity values beyond. (b) Differential propagation phase during the same scan shows clear streaks of attenuation originating in the half-circle of large hydrometeors observed in the reflectivity imagery. (From Palmer et al. 2011, their Fig. 13.)

A second important source of radar product contamination is more organic in nature: large groups of insects and birds can create anomalous reflectivity patterns, known as biological scatter or bioscatter (Gauthreaux 2006). In recent papers, Zhang et al. (2005) and Liu et al. (2005) have laid the groundwork for an automated algorithm aimed at detecting bird migration radar echoes, using a combination of observations and verification via polarimetric radar measurements. Martin and Shapiro (2007) build on this work, and also discriminate between the clear-air echoes caused by birds and insects, primarily by examining radar cross-sections of point targets, thereby determining the target density. While these clear-air returns can be used to obtain information about the atmosphere that would otherwise be invisible to radar imagery (e.g., information about the convective boundary layer (Chandra et al. 2010) or other data in the budding field of radar aeroecology (Chilson et al. 2012)), they also contribute to the non-meteorological contamination of reflectivity data and subsequent quantitative studies. Identification and correction of these returns is thus necessary to ensure the quality of any quantitative analysis of reflectivity and/or velocity.

The focus of this study will be on a third source of error in radar products: ground clutter, where fixed objects near the transmitter block the radar pulse. Contrary to the problem of attenuation identification, a simple map of these fixed ground clutter regions can be created by identifying persistent clear-air echoes in the reflectivity imagery, since the location of patches of ground clutter is more stable than their intensity (Joss 1981). For more precise mapping of clutter that takes into account the temporal characteristics of clutter, the fact that the power spectrum for ground clutter is localized around zero velocity means that identification can be made using Doppler signal processing (Evans and Drury 1983). A scheme to create clutter maps (identifying variable anomalous propagation echoes as well as regular ground clutter) is described by Bellon and Kilambi (1999), using vertical integration of Doppler velocity, as well as the vertical and horizontal gradients of reflectivity. Since the identification of ground clutter is relatively well-understood, especially compared with the different sources of radar product error described above, this study will focus on the correction of data known to be contaminated by stationary ground clutter.

1.3 Ground clutter and the correction of contaminated pixels

For the remainder of this study, the assumption will be that regions of contaminated pixels have been accurately and thoroughly documented using the methods of the previous section. Each pixel contaminated by ground clutter has been identified and mapped. The second part of the problem now applies: how best to correct these contaminated pixels? Can this be done in near-real time, to assist with forecast decision-making, without sacrificing the accuracy necessary for quantitative processing of research data?

A simple method of ground clutter correction is described by Sánchez-Diesma et al. (2001), hereafter SD01, taking an approach consistent with precipitation type. An algorithm is suggested based on two simple schools of thought when it comes to radar contamination correction. The first, advanced by Lee et al. (1995) and Bellon and Kilambi (1999), suggests that a pixel contaminated by ground clutter should simply be replaced with the first non-contaminated pixel directly above it in the vertical. The second, as described by Galli (1984), involves a simple distance-weighted horizontal interpolation over contaminated pixels. Using aspects of both approaches, SD01 suggests a two-step method: first, a thresholding approach is used to differentiate between stratiform and convective precipitation (in SD01, this threshold is taken to be the presence of 45 dBZ reflectivity). Next, the vertical substitution method is applied to convective precipitation and the horizontal interpolation method is applied to stratiform precipitation. The results, reproduced in Table 1.1 below, show that combining horizontal and vertical correction methods provides, on the whole, a better estimation of the true precipitation than using each method individually.

Table 1.1: Comparison of three pixel-substitution methods for instantaneous rainfall rates (mm/hr) of a 24-hour precipitation event featuring a mix of convective and stratiform precipitation. (From Sánchez-Diezma et al. 2001, their Fig. 2.)

	Vertical	**Horizontal**	**Horizontal+Vertical**
Efficiency	0.68	0.73	0.82

r^2	0.72	0.75	0.84
Average Error	0.96	1.02	1.04

Given that stratiform precipitation has strong horizontal homogeneity and that convective precipitation tends to have strong vertical homogeneity, this result is intuitively not surprising. It is worth noting, however, that while straightforward horizontal interpolation performs fairly well on its own, regardless of precipitation type, the addition of the vertical dimension consistently provides an improvement in efficiency and correlation (r^2). Likewise, the ground clutter correction algorithm developed over the following chapters will begin by using a 2D paradigm, evaluating several correction methodologies in the horizontal, and then will gradually add complexity by considering data in the height dimension—and, eventually, in the time dimension as well.

The use of the time dimension in the interpolation of reflectivity data is quite novel: the ready availability of data from different radar scans is a relatively untapped source of information. The reflectivity data from past radar scans alone is an instantly available source of information to help overcome data holes in the spatial dimension, such as widespread stationary areas of ground clutter. If the meteorological situation is not so urgent that immediate, real-time results are necessary, data from future radar scans can also be used. By considering data from five or ten minutes before and after a given scan, a clear picture of storm evolution can be created to help interpolate over areas of missing data. Complementing the spatial information with this fourth dimension suggests, intuitively, that a more accurate rendition of the true reflectivity field can be reconstructed.

1.4 Outline and objectives

This chapter has described the general problem of radar contamination, splitting the issue into two parts. First, how do we identify contaminated pixels? Second, how do we correct pixels we know to be contaminated? The first question is outside the scope of this study, although a brief overview of identification methods for various types of radar contamination was provided in Section 1.2. The focus here will be on correcting a region already known to be contaminated by ground clutter, and will seek to determine how best we can use our knowledge of uncontaminated pixels to make a correction in near-real time.

Chapter 2 introduces the S-band radar data that will be used in this study, describing the ground clutter identification algorithm and additional preprocessing of the raw reflectivity scans. More detailed descriptions are given of the six convective and stratiform events that will form the basis of this study's datasets. Using the data from these events, Chapter 3 describes the single-pixel replacement problem, thus building an understanding of the error structure of the convective and stratiform events in the x-, y-, z-, and t-dimensions. An exploration of the time dimension's error equivalence establishes quantitative evidence of the advantages of blending reflectivity data from multiple times.

With these preliminary single-pixel explorations complete, Chapter 4 lends its focus to a more complex aspect of the ground clutter correction problem: what is the optimal combination of pixels that can be used to replace contaminated data? For this work, reflectivity will be the quantity interpolated. A "smart ordinary kriging" algorithm will be developed, which uses a "bowtie" method to select pixels providing an optimal combination of spatial information while minimizing unwanted smoothing of reflectivity data. Using the datasets described in Chapter 2, as well as the error structure established in Chapter 3, the results of this algorithm are then compared with simpler approaches for clutter correction, such as nearest-neighbour methods and the SD01 approach described in Section 1.3 above. The results will also be examined in the context of different radar products derived from the raw reflectivity values, including instantaneous rainfall rates and one-hour rainfall totals. These algorithms will also be evaluated for high-intensity pixels alone. Finally, Chapter 5 summarizes this study and suggests future work that would allow for the application of this algorithm to other fields, data types (such as Doppler velocity), and radar contamination (such as attenuation).

To summarize, the objectives of this study are as follows:

1 To create an algorithm that can be operated in real time to correct regions of ground clutter in the McGill radar.

2 To build and explore the four-dimensional error structure of reflectivity in convective and stratiform events, in order to quantify the value added by the temporal dimension in radar reflectivity products.

3 To establish the value of blending radar data at different times, including the potential need for regenerating some radar products to take advantage of near-future information.

4 To discuss the implications and future applications of this work to improve a broader set of radar products.

Chapter 2: Description of reflectivity data and preprocessing

2.1 Introduction

In order to begin the process of building an algorithm capable of automatically correcting known contaminated pixels in near-real time, it is necessary to create a comprehensive picture of the error structure of reflectivity. In this study, five convective events and one stratiform event were considered, all of which resulted in substantial precipitation totals over the McGill radar's domain. By splitting some of these events into early and later stages (resulting in a total of six convective and three stratiform "cases"), it was possible to examine how storm growth and decay processes individually affect the reflectivity error structure. The following sections describe these events in some detail, leaving aside synoptic arguments of causality in favour of analysis of the storm structure and geometry, with the aim of creating an automated system capable of replacing contaminated pixels using only knowledge of radar reflectivity data at different times, heights, and locations.

2.2 Preprocessing of radar data

The reflectivity data to be used in this study consist of five convective events and one stratiform event (described in the following section) collected by the McGill S-band radar at the J.S. Marshall Radar Observatory in Sainte-Annede-Bellevue, some thirty km west of downtown Montreal, Quebec, Canada. Figure 2.1 displays the location of the radar. Using the reflectivity data collected at 24 elevation angles (from 0.5° to 34.4°), CAPPIs (constant altitude plan position indicators) are constructed for each five-minute timestep at heights of 1.5 km, 2 km, 2.5 km, 3 km, 3.5 km , 4 km, 5 km, 6 km, 7 km, and 8 km AGL.

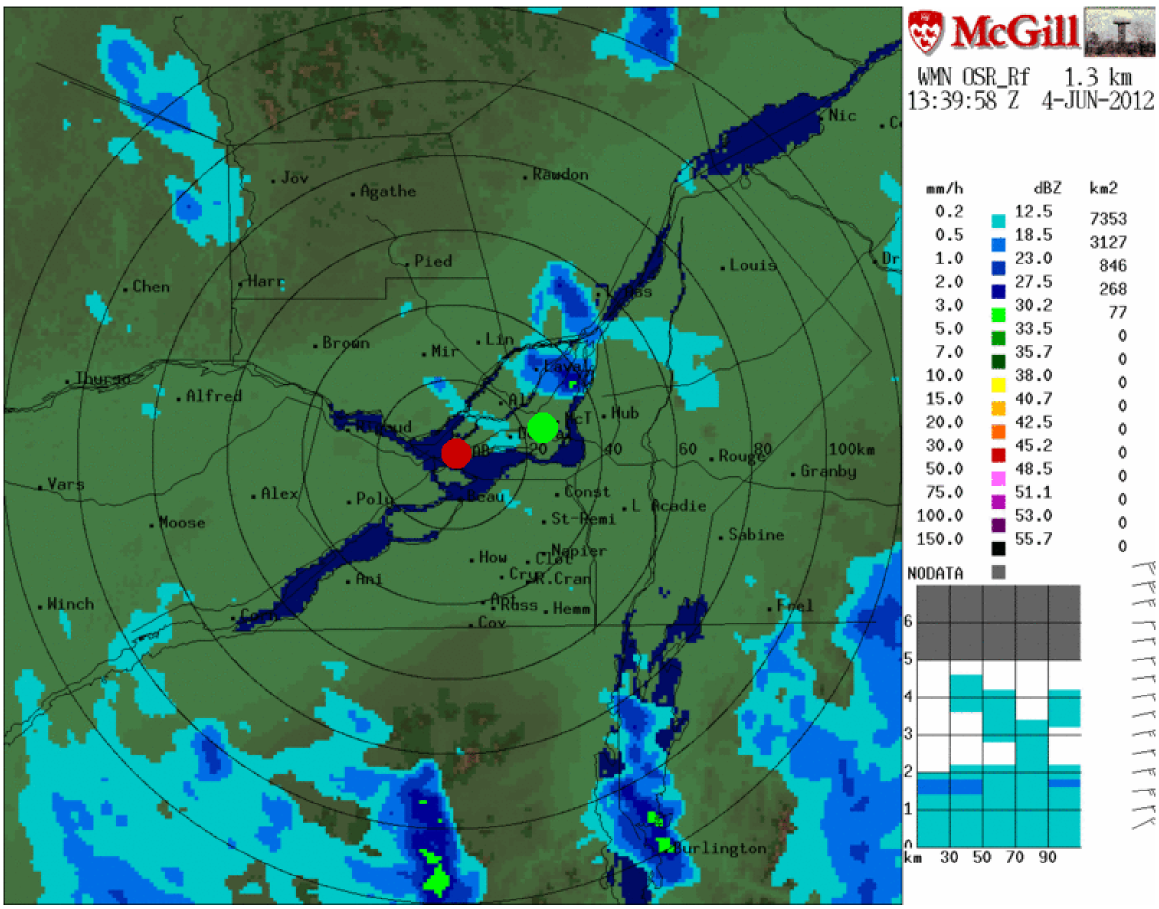

Figure 2.1: Location of the J.S. Marshall Radar Observatory, indicated by the red circle. Downtown Montreal is marked by the green circle. 25-km range rings are spaced around the radar's location. (From http://www.radar.mcgill.ca/imagery/scanning-radar.html.)

Areas of known stationary ground clutter at the McGill radar are established by examining radar returns during days without precipitation. Using archival data, this ground clutter is identified, and a binary array is created for each of the CAPPI levels indicating pixels with known ground clutter (see Fig. 2.2). For the purposes of the single-pixel analysis to be performed in Chapter 3, all of these contaminated pixels were removed from consideration.

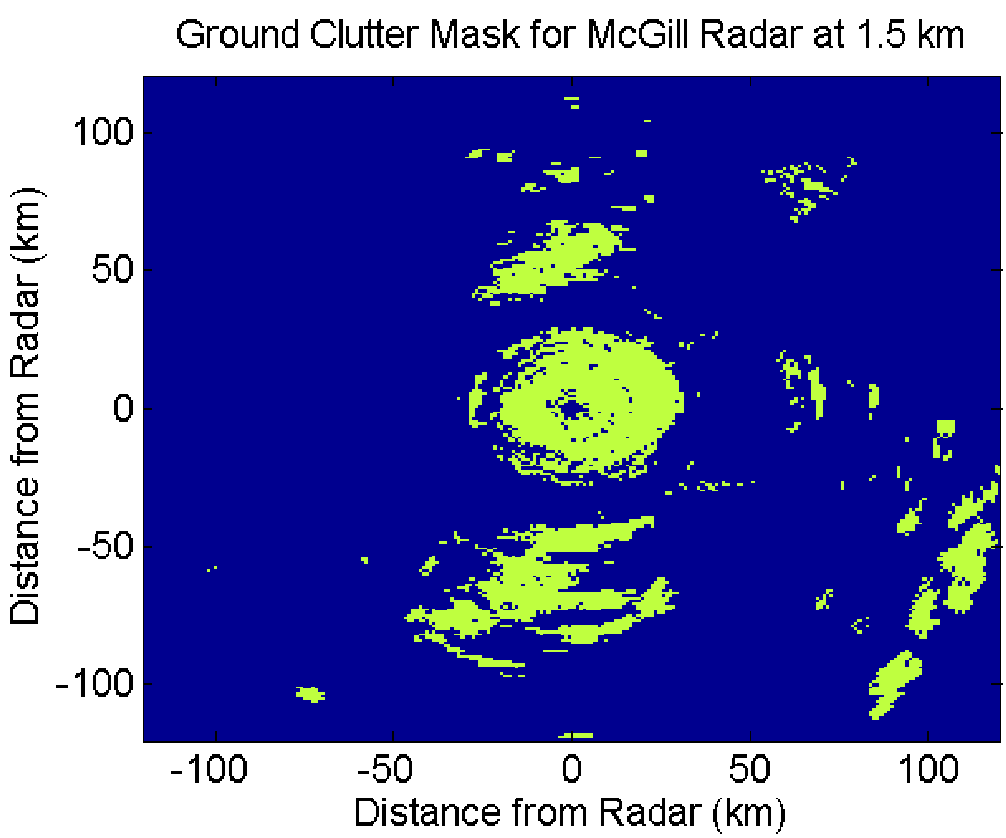

Figure 2.2: Depiction of the known stationary ground clutter at the McGill radar for the 1.5 km CAPPI height. Locations of pixels contaminated by ground clutter

are indicated in green, and each pixel represents a 1 km × 1 km horizontal square. The domain is 240 km × 240 km.

2.3 Description of precipitation events

Six rainfall events were selected (five convective and one stratiform) based on their intensity and variety of meteorological features. It is worth noting that, rather than using an explicitly threshold-based distinction between convective and stratiform precipitation (as in Bellon and Kilambi, 1999), the events here have been classified based on a more general overall observation of the precipitation type. This has resulted in a series of events that are very clearly either stratiform or convective, with little overlap between the two. Figure 2.3 depicts simple snapshots of these events as they appear after the preprocessing described in the previous section. A detailed, causal description of synoptic and mesoscale forcings for each event is not included, as the ground clutter correction algorithm built over the following chapters derives its information from radar scans alone; knowledge of the geometry of the precipitation echoes is sufficient. Table 2.1 provides a summary of the six events.

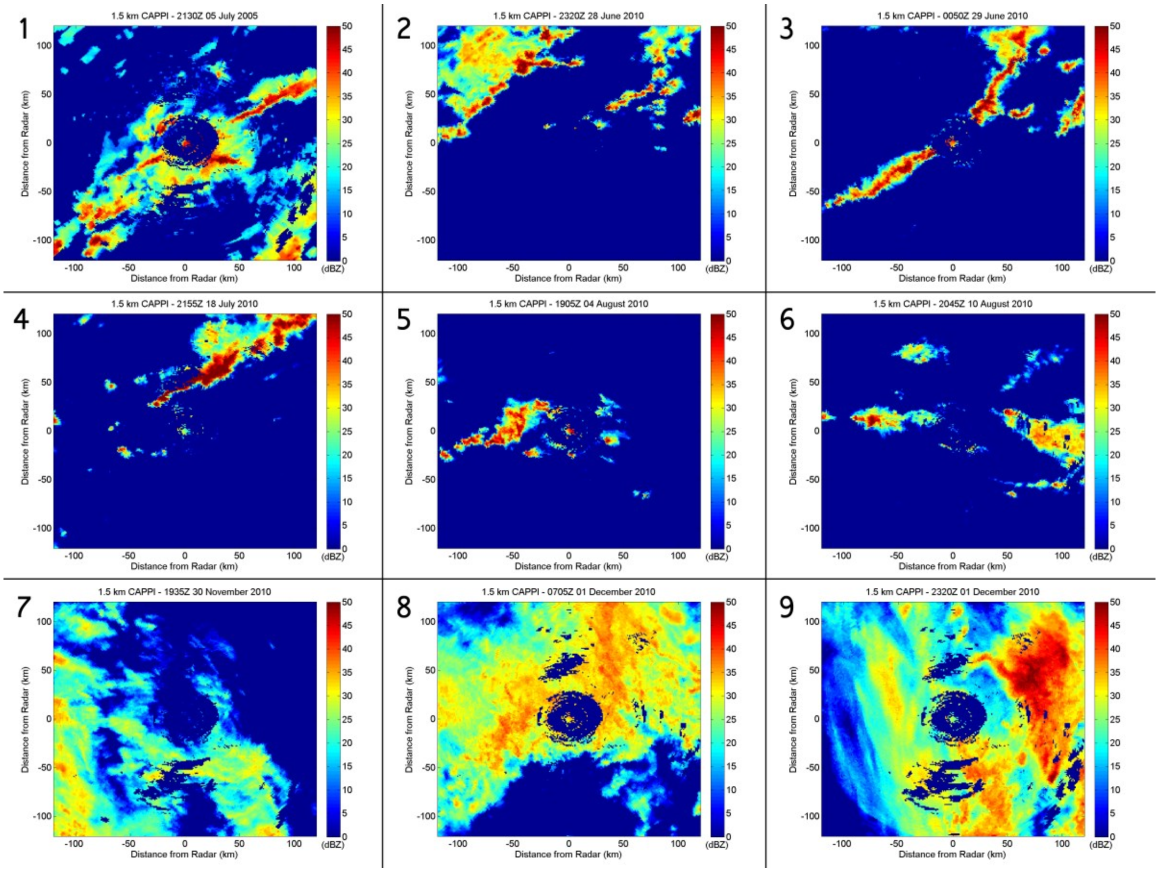

Figure 2.3: Sample reflectivity scans for each of the nine events considered in this study. The preprocessing described in Section 2.2 has been performed, as evidenced by the regions of blotted-out ground clutter corresponding with those in Fig. 2.2. See text for a description of the events. Case numbers (top left of each part of the figure) correspond with those used in Table 2.1.

Table 2.1: Summary of events considered during this study. A convective event is characterized by organized cumuliform precipitation with strong vertical development, while a stratiform event features moderate precipitation with substantial horizontal extent.

Case	Convective or Stratiform?	Start Time	End Time
Case 1	C	1905Z 05 Jul 2005	2200Z 05 Jul 2005
Case 2	C	2005Z 28 Jun 2010	0000Z 29 Jun 2010
Case 3*	C	0005Z 29 Jun 2010	0200Z 29 Jun 2010
Case 4	C	1700Z 17 Jul 2010	0000Z 18 Jul 2010
Case 5	C	1700Z 04 Aug 2010	2200Z 04 Aug 2010
Case 6	C	1700Z 10 Aug 2010	2200Z 10 Aug 2010
Case 7	S	1905Z 30 Nov 2010	0000Z 1 Dec 2010
Case 8*	S	0005Z 01 Dec 2010	1200Z 01 Dec 2010
Case 9*	S	1200Z 01 Dec 2010	0000Z 02 Dec 2010

* see text for discussion

The six rainfall events were split into 9 "cases"—for instance, the rainfall event on 28-29 June 2010 was split into two cases, so that the error structures described in the following chapter could be computed and compared at both the early and late stages of the same event. The same approach was taken for the single, long-lasting stratiform event, which was split into three cases. Given the computational resources necessary to process such a long-lived event, only one stratiform rainfall event was included in this dataset. In addition, to reduce complexity due to the ice phase and the bright band, only liquid rainfall events were considered.

As mentioned above, the events chosen for this study were selected on the basis of the severity of their rainfall (the stratiform event in cases 7-9 resulted in 45 mm of precipitation, and the convective events of Case 1 and Case 5 contributed to 48-hour precipitation totals of 50 mm and over 100 mm, respectively), as well as their appearance as "typical" convective or stratiform storms for the region. Common features among the six convective cases are apparent, looking at the first two rows of Fig. 2.3: in the domain covered by the radar imagery, convection frequently organizes into a line oriented either SW-NE or W-E, and propagation of these systems has a westerly component, though new cell development is nearly always along the western edge of the storm. The following chapter will determine how far these common features extend into the statistical error structure of each event.

These nine cases provide an interesting variety of weather events that have had important impacts on the public: by comparing the statistics of these events, a better picture of the overall variability and magnitude of reflectivity error structure can be constructed. The following chapter approaches this endeavour by considering the single-pixel replacement problem, paving the way for more complex approaches in Chapter 4.

Chapter 3: The single-pixel replacement problem

3.1 Introduction

With radar data that have been preprocessed and organized, it is now possible to begin an examination of the spatial and temporal error structure of the reflectivity measured in each of the nine cases described above. The next section of this chapter, Section 3.2, describes an important correction that must be made in the context of the replacement of pixels with height. Section 3.3 outlines a methodology for displaying and analyzing reflectivity error structure, depicting and comparing examples of this output based on the substantial precipitation events described in the previous chapter. By establishing these quantitative markers based on error structure, we are gaining an understanding of the answer to an important question in the ground clutter replacement process: when is it better to choose our replacement data from different times or heights rather than from the same radar scan? Section 3.4 seeks to answer this question by establishing simple metrics of error equivalence between time, height, and horizontal distance. These same metrics are used in Chapter 4 to aid the final ground clutter replacement algorithm in selecting the best possible combination of pixels to replace contaminated ones.

3.2 Vertical profile of reflectivity

In order to make use of radar data at different heights, one additional routine preprocessing correction must be made, involving the vertical profile of reflectivity (VPR, see Koistinen 1991). For stratiform events in general, the problem of the radar bright band in the melting layer is a substantial one: the error structure of single-pixel reflectivity replacement would be skewed by these unrealistic returns. To correct for this effect, among others, a simple vertical profile of reflectivity is constructed for each radar scan by calculating the mean reflectivity at each CAPPI height level over the 20 minutes preceding and succeeding the scan in question. This profile is then subtracted from the reflectivity data at each height, creating a dataset that allows more accurate reflectivity intercomparisons from one height to another. Figure 3.1 shows an example of a VPR used in the correction process. While this VPR correction is simpler than many correction methods in use today (Bordoy et al. (2010) include a cautionary note warning that the VPR's rapid variation in time and space makes correction a difficult prospect), it was deemed sufficient for the purposes of this study.

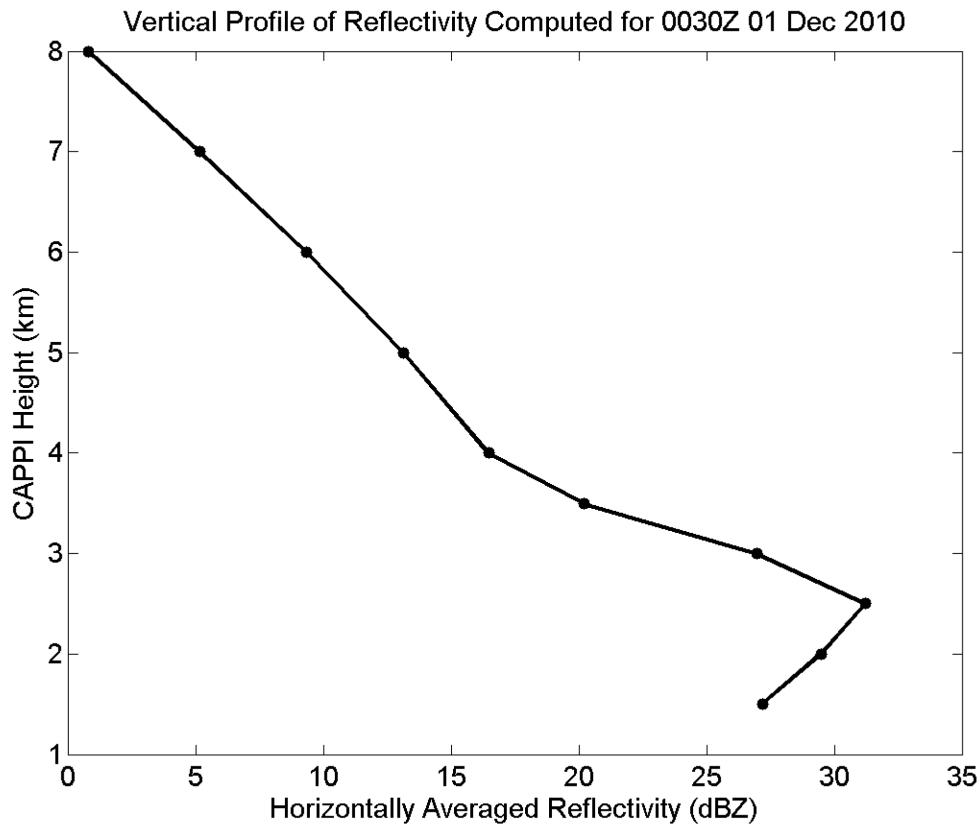

Figure 3.1: VPR correction used for 0030Z on 01 December, 2010. This stratiform precipitation is contained in Case 8 (see discussion in Section 2.3). Values (in dBZ) represent the height- and time-averaged reflectivity measured at each of the CAPPI heights over a forty-minute period centred on the time in question.

3.3 Four-dimensional error structures

Before tackling the problem of replacing large areas of ground clutter (Fig. 2.2) with the best possible combination of surrounding data in time and space, it is necessary to consider a more simplified version of the issue at hand. Consider the problem of a single contaminated pixel which we wish to replace with a single pixel from elsewhere in the spatio-temporal domain. How much error will result for each possible pixel replacement? Since reflectivity is essentially a continuous spatial variable, we expect a gradual drop-off in accuracy as the distance between our replacement pixel and the contaminated pixel increases. One of the key motivations of this study, however, involves quantifying the benefit of considering data from different times, as well as in space. In what situations is a pixel from five minutes earlier going to be a better match than a pixel from the same radar scan? By gaining a detailed knowledge of this four-dimensional error structure (hereafter referred to as a variogram), it becomes possible to create more complex algorithms that can select replacement pixels based on minimizing statistical error as much as possible.

In order to examine various types of precipitation, the nine cases listed in Table 2.1 are considered separately, and statistics are computed for each. The procedure is a simple one, if computationally expensive. For each case:

1 A reference pixel $Z(x_0, y_0, z_0, t_0)$ is selected.

2 This reference pixel exists on a single radar scan in the x-y plane (such as a CAPPI or a PPI), i.e., the plane of pixels having values $Z(x, y, z_0, t_0)$, where $x = x_0 + \Delta x$ and $y = y_0 + \Delta y$. The squared difference between the reference pixel and its surroundings is calculated for each pixel contained in this plane,

resulting in a 2D matrix containing values

$$\sigma^2(\Delta x, \Delta y, z_0, t_0) = [Z(x_0, y_0, z_0, t_0) - Z(x, y, z_0, t_0)]^2 .$$

3 Similar matrices of squared differences can be built for the scans at the other 9 CAPPI height levels, i.e., the planes of pixels having values $Z(x, y, z, t_0)$, with x and y defined as above, and $z = z_0 + \Delta z$. The result is a series of ten 2D matrices $\sigma^2(\Delta x, \Delta y, \Delta z, t_0) = [Z(x_0, y_0, z_0, t_0) - Z(x, y, z, t_0)]^2$. Note that, as mentioned in the previous chapter, any pixels contaminated by ground clutter have been omitted from these and future computations.

4 Finally, we consider the scans at earlier and later times (established somewhat arbitrarily as ±25 minutes, making for a total of eleven timesteps), i.e., the planes of pixels having values $Z(x, y, z, t)$, with x, y, and z defined as above, and $t = t_0 + \Delta t$. The result is a series of eleven sets of ten 2D matrices

$$\sigma^2(\Delta x, \Delta y, \Delta z, \Delta t) = [Z(x_0, y_0, z_0, t_0) - Z(x, y, z, t)]^2 .$$

5 We have now established a four-dimensional variogram (a description of the error structure in (x, y, z, t) for a single reference pixel). The next step is to generalize this across an entire event by performing steps one through four for every possible reference pixel at every possible time. For Case 1, which consists of 36 timesteps, each containing a 240 × 240 grid of reflectivity values for each of the ten CAPPI heights, the final result will be 18,432,000

sets of 110 2D squared-difference matrices $\sigma^2(x, y, z, t)$--a substantial store of information.

6 The final step before analysis of the data is possible is to average all of these sets of $\sigma^2(x, y, z, t)$, resulting in a representative four-dimensional variogram (in the form of 2D matrices of variance for each of the ten height levels at each of the eleven timesteps) that describes the error structure for the entire event.

3.3.1 Convective events

An illustration of the end result of this process is in Fig. 3.2, which shows the 2D variogram for Case 1 at a given time and height—in other words, it shows the x-y error distribution for single-pixel replacement, not considering variations in height or time.

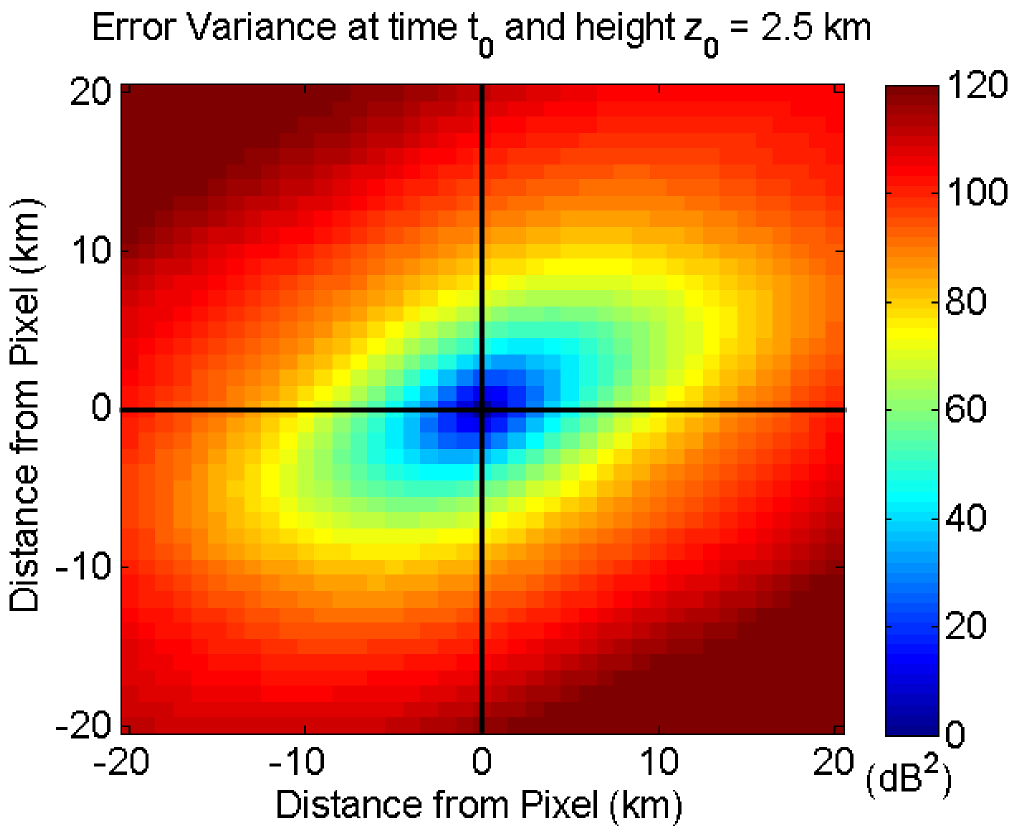

Figure 3.2: 2D variogram calculated for Case 1 at time t_0 and height z_0 = 2.5 km. This image can be thought of as a way to estimate how much error would result if we were to replace a reference pixel with a pixel from its surroundings. The crosshairs are centred on the location of the reference pixel. In terms of distance, 1 pixel = 1 km. Since the variance involves the sum of a difference in reflectivities [dBZ], the units of this figure are dB^2.

Note that the SW-NE oriented axis visible in Fig. 3.2 corresponds with the orientation of the axis of the linear convective feature that dominates in this particular weather event (recall Fig. 2.3). Essentially, the error structure confirms that pixels located along the SW-NE axis are more likely to be similar (i.e., to have lower values of variance) than pixels along the perpendicular NW-SE axis. This tendency is, of course, due to the along-line homogeneity of a squall line.

Given that all six convective cases featured squall lines oriented in roughly the same direction (a common orientation for organized convection in the Montreal area), it is expected that the shape of the variogram will be roughly similar for all six cases. Figure 3.3 compares 2D variograms for all six convective cases, where Fig. 3.3a is identical to Fig. 3.2, and the remaining parts of the figure are also taken at time t_0 and height 2.5 km. The shapes of the variograms are again reflective of the general orientation of the storm structure. For instance, when comparing the shapes of Figs. 3.3b and 3.3c, we might expect that the latter case will feature a more strongly linear feature, which is indeed true (cf. Fig. 2.3). Recall also that Cases 2 and 3 are two halves of the same event (cf. Table 2.1): as the line begins to break down late in the event (Case 3: Fig. 3.3c), we see a variogram with an elliptical shape that is less stretched along the orientation of the precipitation line.

The elliptical shape of the variogram for Fig. 3.3e shows much more of a W-E orientation in terms of its longer axis, which reflects the W-E oriented line of convection that dominates the period (cf. Fig. 2.3). Likewise, while the dominant convective feature in Case 6 was much more W-E oriented than the tilted slant of

Fig. 3.3f, it is worth noting that later in the period, a thunderstorm outflow boundary forced new convection in the southeastern quadrant of the domain, which would account for the slant in the variogram's shape.

Another point of interest in Fig. 3.3 is that some events had considerably higher values of variance in their 41 × 41 km variogram domain than others—note the different colour scales used for Figs. 3.3a and 3.3b. For Case 1 in particular (Fig. 3.3a), the variance in the across-line direction does not drop off nearly as quickly as it does in any of the other cases. Looking at the precipitation pattern of this case in comparison with the other convective events, it is apparent that while Case 1's reflectivity values are similar to those in the other events, it does feature much more widespread precipitation than any of the others. Thus, in a squall line, we find that a less sharp precipitation gradient will correspond with lower values of variance in the across-line direction.

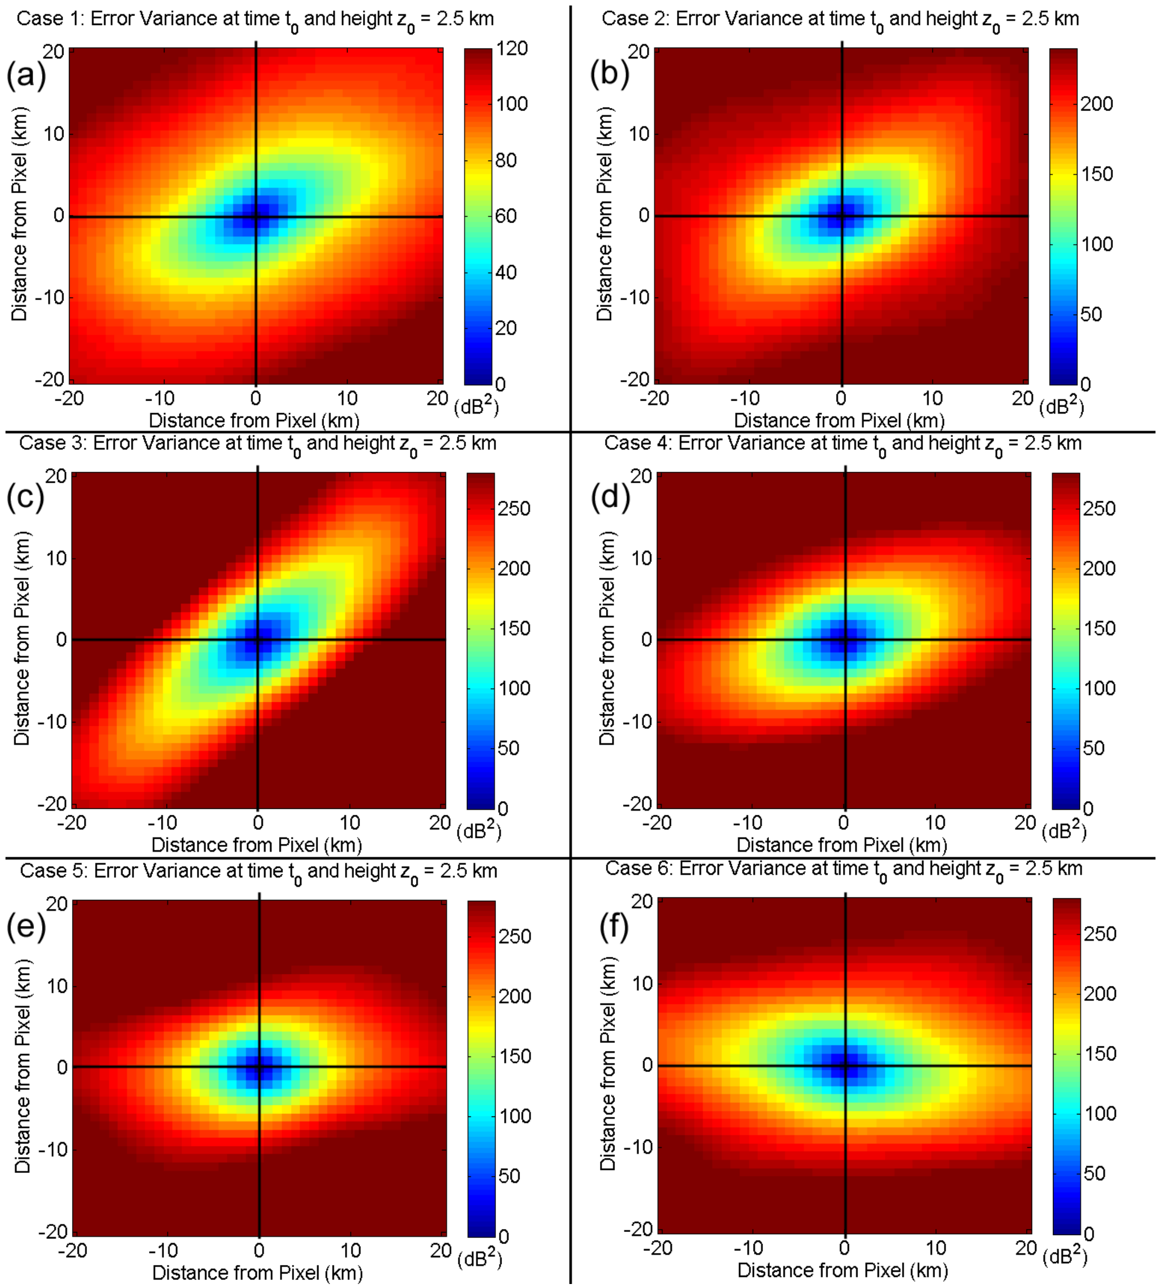

Figure 3.3: 2D variograms calculated for the six convective cases of this study. Note the different colour scales for certain cases: Case 1 showed lower variances overall than any of the other five events. Each figure depicts the variogram at time t_0 and at height 2.5 km (as in Fig. 3.2).

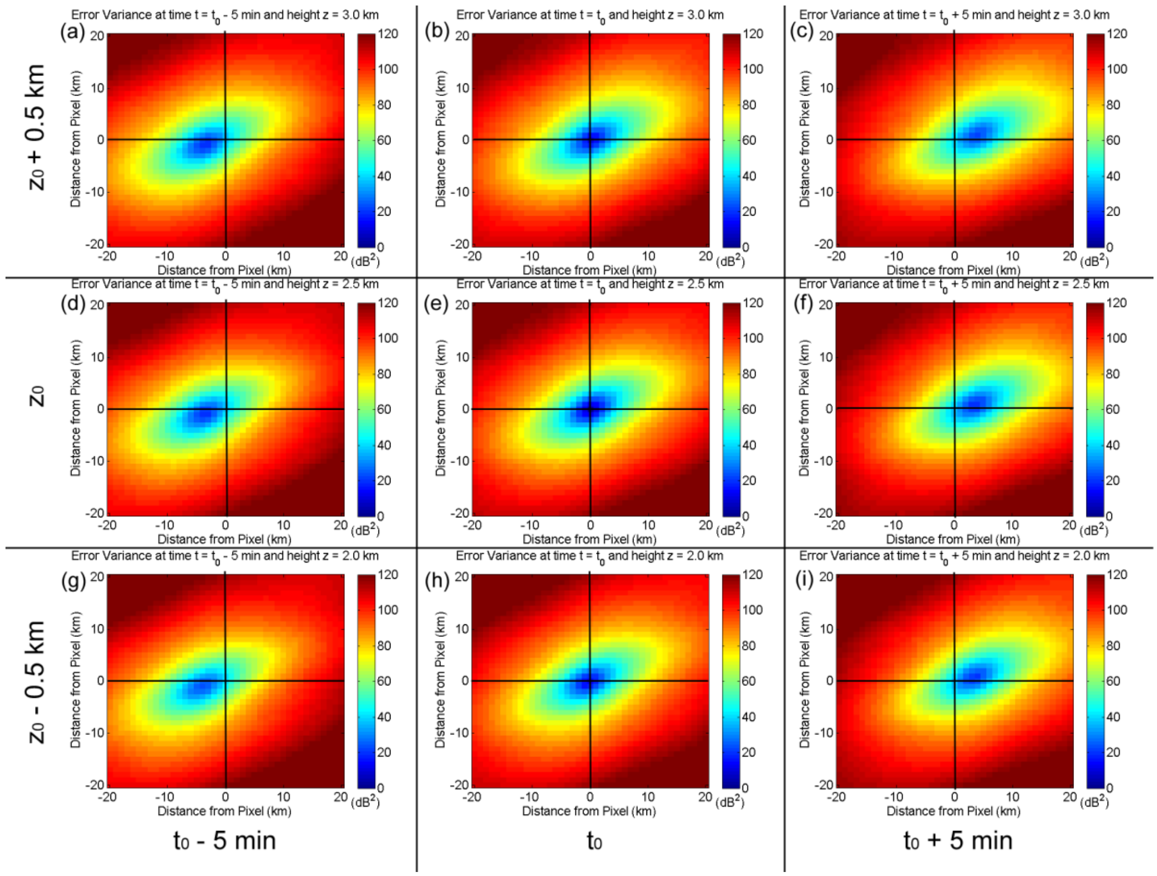

Figure 3.4: 2D variograms calculated for Case 1 as described above. This image can be thought of as a way to estimate how much error would result if we were to replace a reference pixel with a pixel from its surroundings in x, y, z, and t. The centre row is at height $z = z_0$ (in this case, 2.5 km), the top row is at height $z = z_0 + 0.5km$, the bottom row is at height $z = z_0 - 0.5km$, the centre column is at time $t = t_0$, the left column is at time $t = t_0 - 5\min$, and the right column is at time $t = t_0 + 5\min$.

As discussed in Chapter 1, when correcting for ground clutter in reflectivity data, it is limiting to consider only information in the horizontal—an understanding of the variability in height and in time also provides valuable information. Figure 3.4 shows a subset of the Case 1 variograms centred at 2.5 km height. The middle of the three rows represents 2.5 km height (defined as $z = z_0$), the upper represents 3 km height, and the lower represents 2 km height. The middle column shows variance information for time $t = t_0$ (i.e., the time at the radar scan containing gaps to be filled), the left column shows variance information for time $t = t_0 - 5\,\mathrm{min}$, and the right column shows variance information for time $t = t_0 + 5\,\mathrm{min}$. Thus, the centre image (Fig. 3.4e) is the same variogram as that shown in Fig. 3.3. As is expected of a strong convective event, there is a great deal of homogeneity in the vertical over the 1-km range under consideration—given a region of ground clutter, this image clearly shows that there are excellent candidates for replacement pixels at nearby heights. Likewise, the variograms at earlier and later times contain a great deal of useful data (i.e., low variance)—the benefit of using data at different times is visible in the error structure depicted there. We can also see the propagation of the storm toward the ENE by following any of the individual rows of the figure.

3.3.2 Stratiform events

Having examined the relevant variograms for the convective events, we turn our attention to the stratiform cases. Figures 3.5a through 3.5c show the variograms at time t_0 and height 2.5 km for Cases 7-9, respectively. Recall that these three cases represent the early, middle, and final portions of the same 29-hour stratiform precipitation event (cf. Table 2.1). In the early stages of the event (Fig. 3.5a), we see an elliptical shape with its longer axis oriented NW-SE. This makes sense—for the first five hours of the event, the precipitation is still propagating into the radar's domain in the form of bands oriented NW-SE. Later in the event (Figs. 3.5b and 3.5c), we see a much more symmetrical variogram, indicating more evenly spread precipitation, which is indeed what was observed in Cases 8 and 9.

Another feature of interest is that, noting that the same colour scale was used for Figs. 3.5a through 3.5c, there are generally higher variances present in the later stages of the storm (i.e., Fig. 3.5c). Based on our analysis of the convective events, this stronger variability indicates sharper boundaries between regions of precipitation and regions without precipitation. Late in the event's timeframe, bands of heavier rainfall are propagating through the weaker precipitation, providing the sharper gradients in reflectivity necessary to produce the higher variances in Fig. 3.5c.

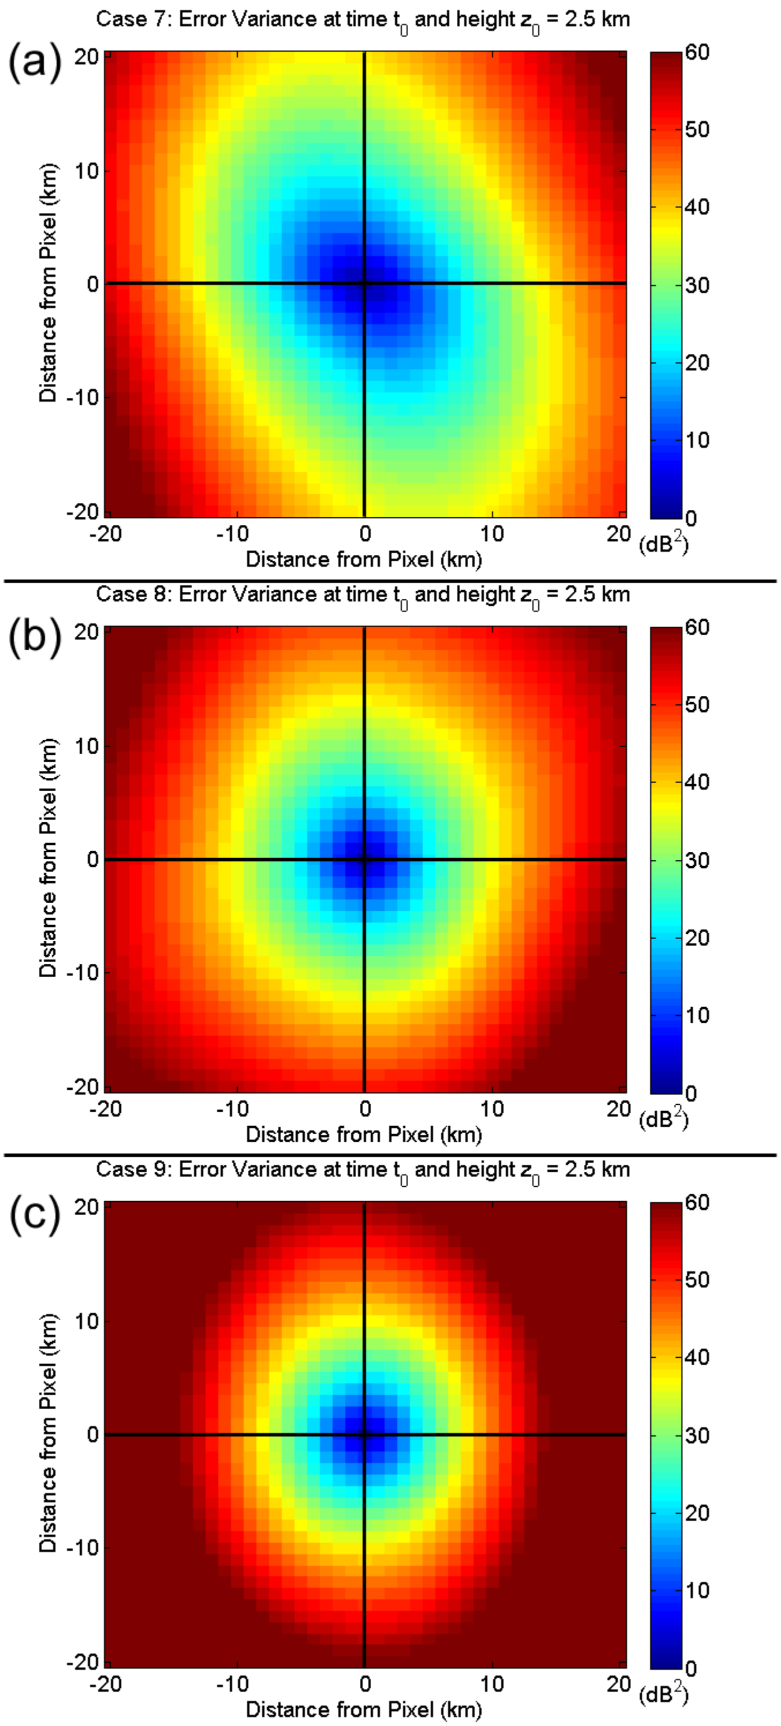
Case 7: Error Variance at time t_0 and height z_0 = 2.5 km
(a)
Distance from Pixel (km)
Distance from Pixel (km)
(dB^2)
Case 8: Error Variance at time t_0 and height z_0 = 2.5 km
(b)
Distance from Pixel (km)
Distance from Pixel (km)
(dB^2)
Case 9: Error Variance at time t_0 and height z_0 = 2.5 km
(c)
Distance from Pixel (km)
Distance from Pixel (km)
(dB^2)

Figure 3.5: 2D variograms calculated for Cases 7-9. Since these cases represent three stages of the same stratiform precipitation event (the first five hours, the next twelve hours, and the final twelve hours, respectively), this image provides an opportunity to compare the early, mature, and dissipating stages of a large area of substantial stratiform precipitation. Note that, for all three images, the same colour scale was used.

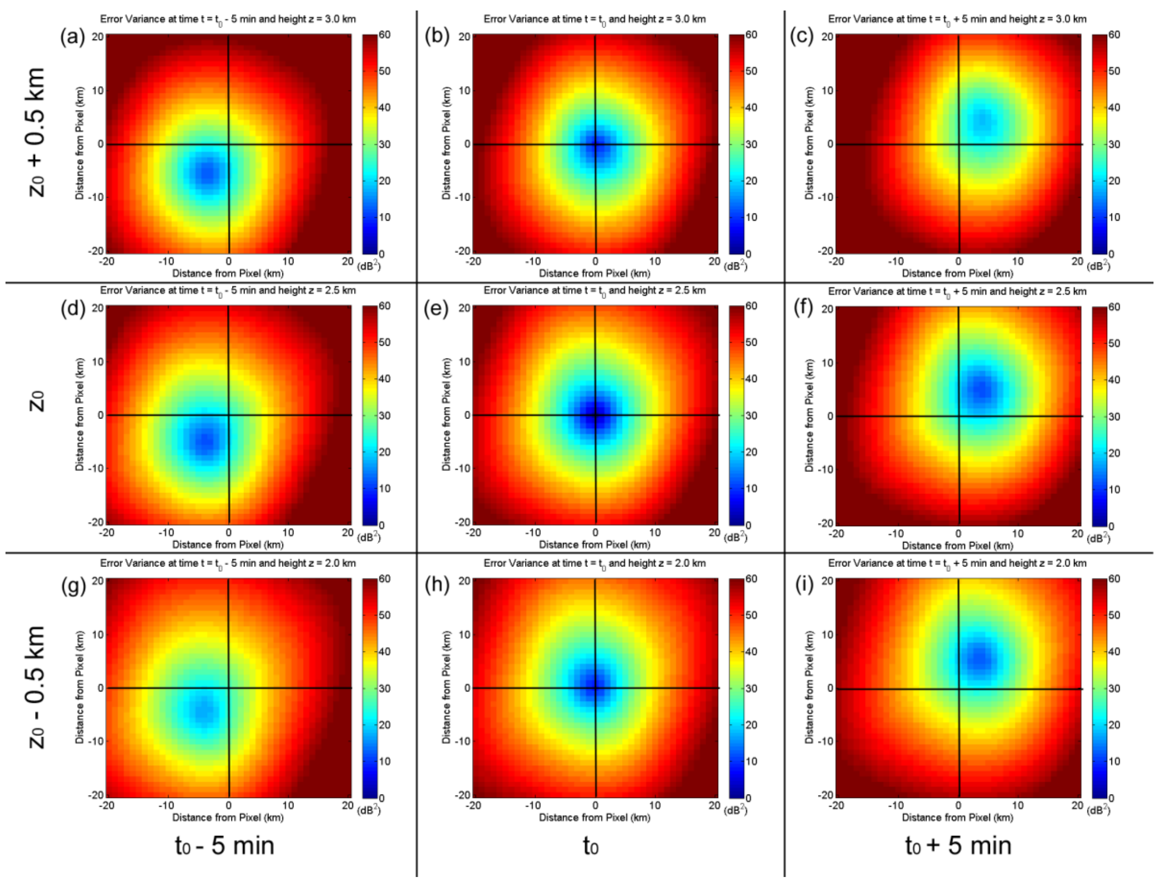

Figure 3.6: As in Fig. 3.4, but for the stratiform Case 8 rather than the convective Case 1. The data for this case come from twelve hours in the midst of a single precipitation event.

Similar to Fig. 3.4, Fig. 3.6 shows the variograms for a single case (in this example, Case 8, the middle times of the stratiform event) at different heights and times. The direction of propagation of the storm is visible across any of the three rows: the precipitation echoes are, on average, propagating from southwest to northeast during this part of the event. Once again, we see a benefit to using data from different times, in the form of the low variances available five minutes before and after the radar scan in question if we take storm displacement into account. Likewise, there is also important information available at different heights, though it appears that this information is less valuable for a stratiform event than it is for a convective event. For instance, if we imagine a situation where we want to replace a pixel at 2.5 km with the best possible pixel at 3.5 km, the variance for a convective event is 15.20 dB^2, while the variance for a stratiform event is 21.49 dB^2. If we instead attempt to replace that pixel with one at 5 km, the difference is still more pronounced: for a convective event, the variance is 38.36 dB^2, while the variance for the stratiform event is 72.54 dB^2.

3.3.3 Comparisons and results

Examining this last point more closely, in Fig. 3.7 we see a height profile of variograms for a convective event (left column) and a stratiform event (right column). As is expected given the more substantial vertical extent of convective precipitation, variance does not drop off as quickly with height as it does for the stratiform event: thus, for convective events we would expect that there is more valuable information at different heights than there is for stratiform events. Conversely, stratiform events show a much slower drop-off in variance with horizontal distance (compare Figs. 3.7c and 3.7f, for example), which reflects the horizontal homogeneity that characterizes stratiform precipitation. There is thus more information to be gained in the x-y plane for a stratiform event than for a convective event. Any algorithm meant to replace ground clutter with a combination of surrounding pixels must take into account this fundamental difference between the variability of these two types of precipitation, as was demonstrated by Sánchez-Diezma et al. (2001).

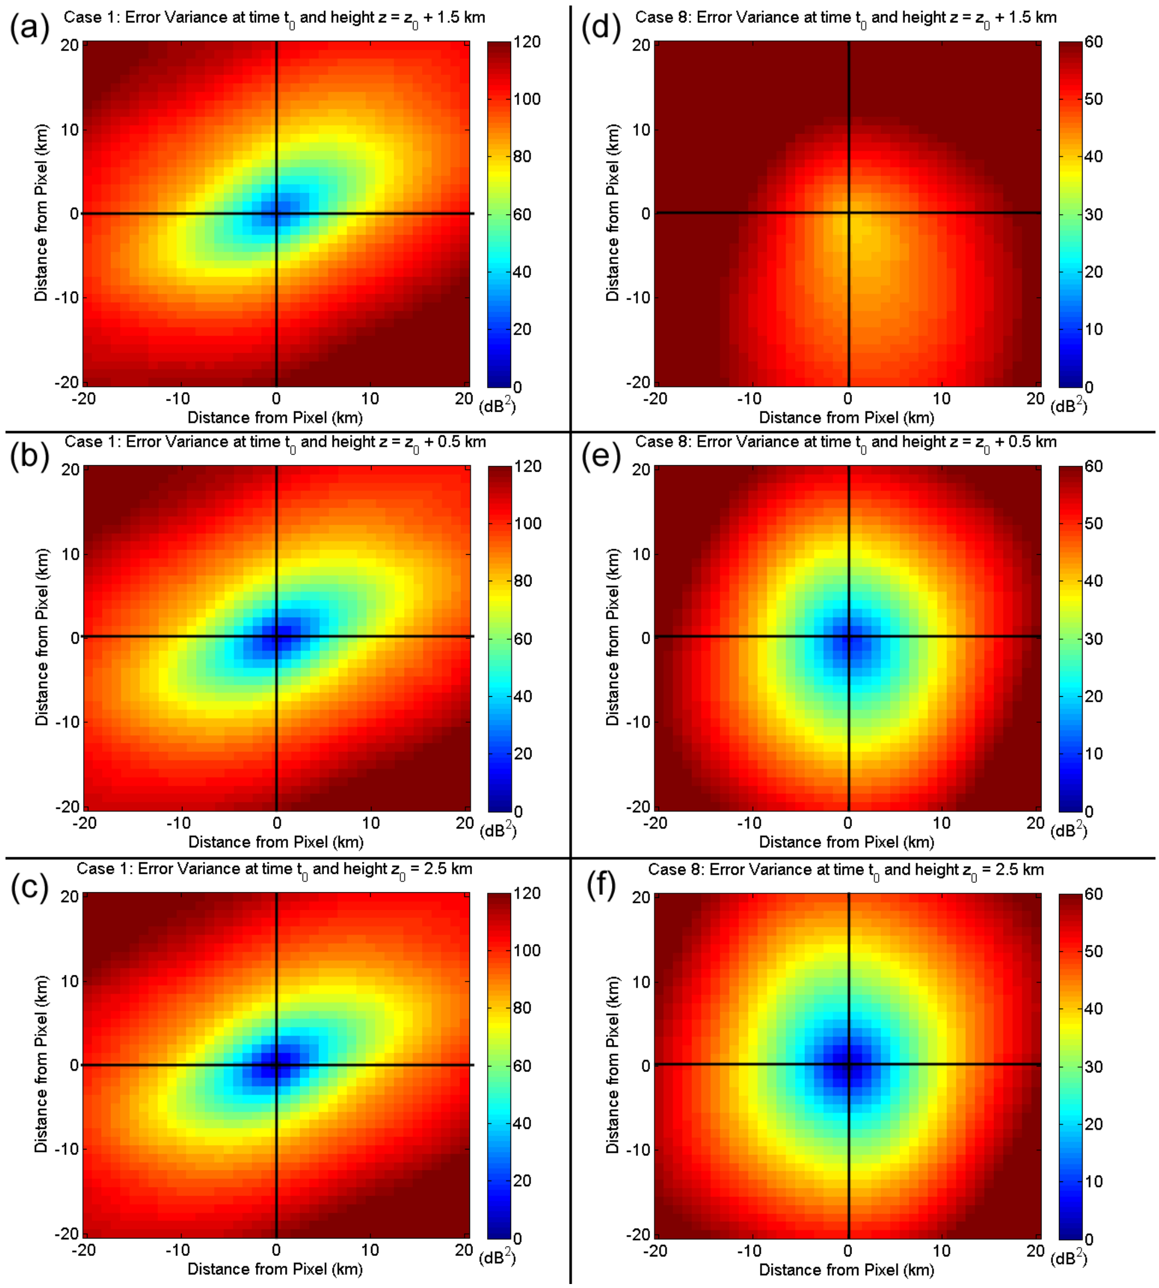

Figure 3.7: Display of variance associated with replacing a pixel at 2.5 km height, comparing spatial variance for a convective case (a-c) and a stratiform case (d-f). Note that the convective case's colour scale is set to double that of the stratiform case—this is to remain consistent with the scales used in previous figures. While the convective case's pattern of variance remains very similar over the 1.5 km

vertical extent of these depictions, the stratiform case's variance drops off rapidly with height. These variance characteristics reflect the more substantial vertical extent of convection.

Figure 3.8 again compares the behaviour of convective and stratiform events, but over time instead of height. Note that, despite the differences in magnitude and horizontal extent of the variance, both the convective and stratiform variograms show a marked similarity between the structure at our reference time $t = t_0$ and at earlier times. This demonstrates another strength of considering data from different times: both stratiform and convective rainfall events have the potential to benefit from this information, since both involve the propagation of organized systems. The lifetimes of convective and stratiform events may be very different (one hour as compared with one day), but we are considering data from a relatively short span of time (under half an hour) and a relatively small spatial scale (a few tens of km), so this difference is less important than it might be otherwise. Thus, while the variability of spatial data (horizontal or vertical) depends strongly on precipitation type, the variability of temporal data relies only on the direction of propagation of precipitation, as is generally determined by the mean echo displacement.

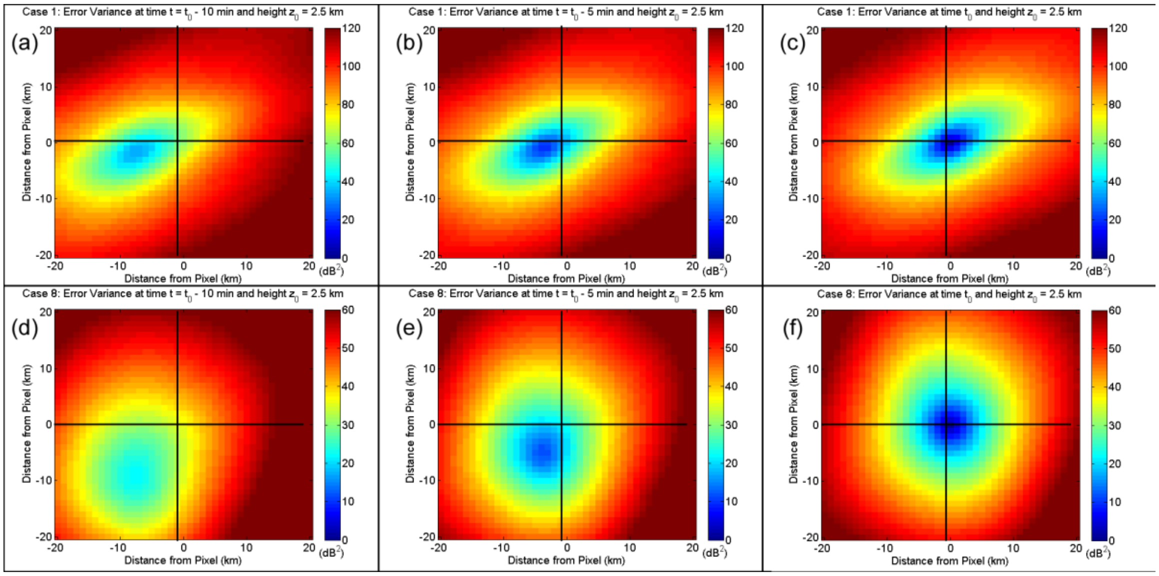

Figure 3.8: Display of variance with time, comparing spatial variance for a convective case (a-c) and a stratiform case (d-f). Note that the convective case's colour scale is set to double that of the stratiform case—this is to remain consistent with the scales used in previous figures. Note that both cases have a similar relative drop-off of variance at earlier times.

With an understanding of the four-dimensional error structure for all nine cases, as well as the points of commonality and differences between convective and stratiform events, we have a solid basis for a system of correction of ground clutter: we know where to look, spatially and temporally, to find information that will be similar to the data contaminated by the ground clutter. The next section takes a more explicitly quantitative approach to the problem through calculations to determine distance-time and distance-height equivalents (i.e., how far do we have to travel horizontally to find a similar variance to the variance five minutes earlier or the variance 1 km higher?). In Chapter 4, an algorithm for ground clutter correction is built using a combination of these equivalence factors and the error structure derived in this section.

3.4 Error equivalence in time, height, and space

The notion of error equivalence provides an important step towards quantifying the benefits of including reflectivity data from different times and different heights in our stable of possible replacement pixels. If we have a pixel contaminated by ground clutter, and the nearest available non-cluttered pixel in the x-y plane is 4 km away, are we better off taking data from 5 minutes earlier? 0.5 km higher? Some combination of the two? A simple metric of error equivalence (more precisely, variance equivalence) will also be essential for the ground clutter correction algorithm outlined in Chapter 4: without knowledge of the relative variance drop-off rates in time and in space, it becomes impossible to automate selection of pixels in all four dimensions. This section establishes the requisite common reference frame that allows comparisons among the four dimensions—it will be possible to calculate relative "distances" in time for use in the kriging approach of the following chapter.

First, we must establish a common metric for height and horizontal distance. As mentioned in Section 2.2, there are ten discrete height levels at which CAPPIs are taken in this dataset (1.5 km, 2 km, 2.5 km, 3 km, 3.5 km, 4 km, 5 km, 6 km, 7 km, and 8 km). For each of these height levels, our goal is to establish an equivalent horizontal distance—at what horizontal distance from our contaminated pixel will a pixel from a different height be a better replacement (i.e., have a lower variance)? In other words, when searching for replacements for contaminated pixels, when should we stop looking in the horizontal and start looking in the vertical?

The process is a simple one. First, we establish a reference height z_0—in keeping with previous figures and examples, we will choose z_0 = 2.5 km as an example. Next, we look for the lowest possible variance at a different height (say, $z = z_0 + 1.5$ km, i.e., $z = 4$ km)—this will be the "bullseye" visible in the variograms of the previous section. Now that we have this value for the variance, we search for it back at our reference height z_0. The horizontal distance at which we find this variance is thus the equivalent horizontal distance for a 1.5 km rise in height. That is, if we travel further than this horizontal distance in search of a replacement pixel, we should instead consider looking at a CAPPI level 1.5 km higher up. Figure 3.9 illustrates this process for Case 1.

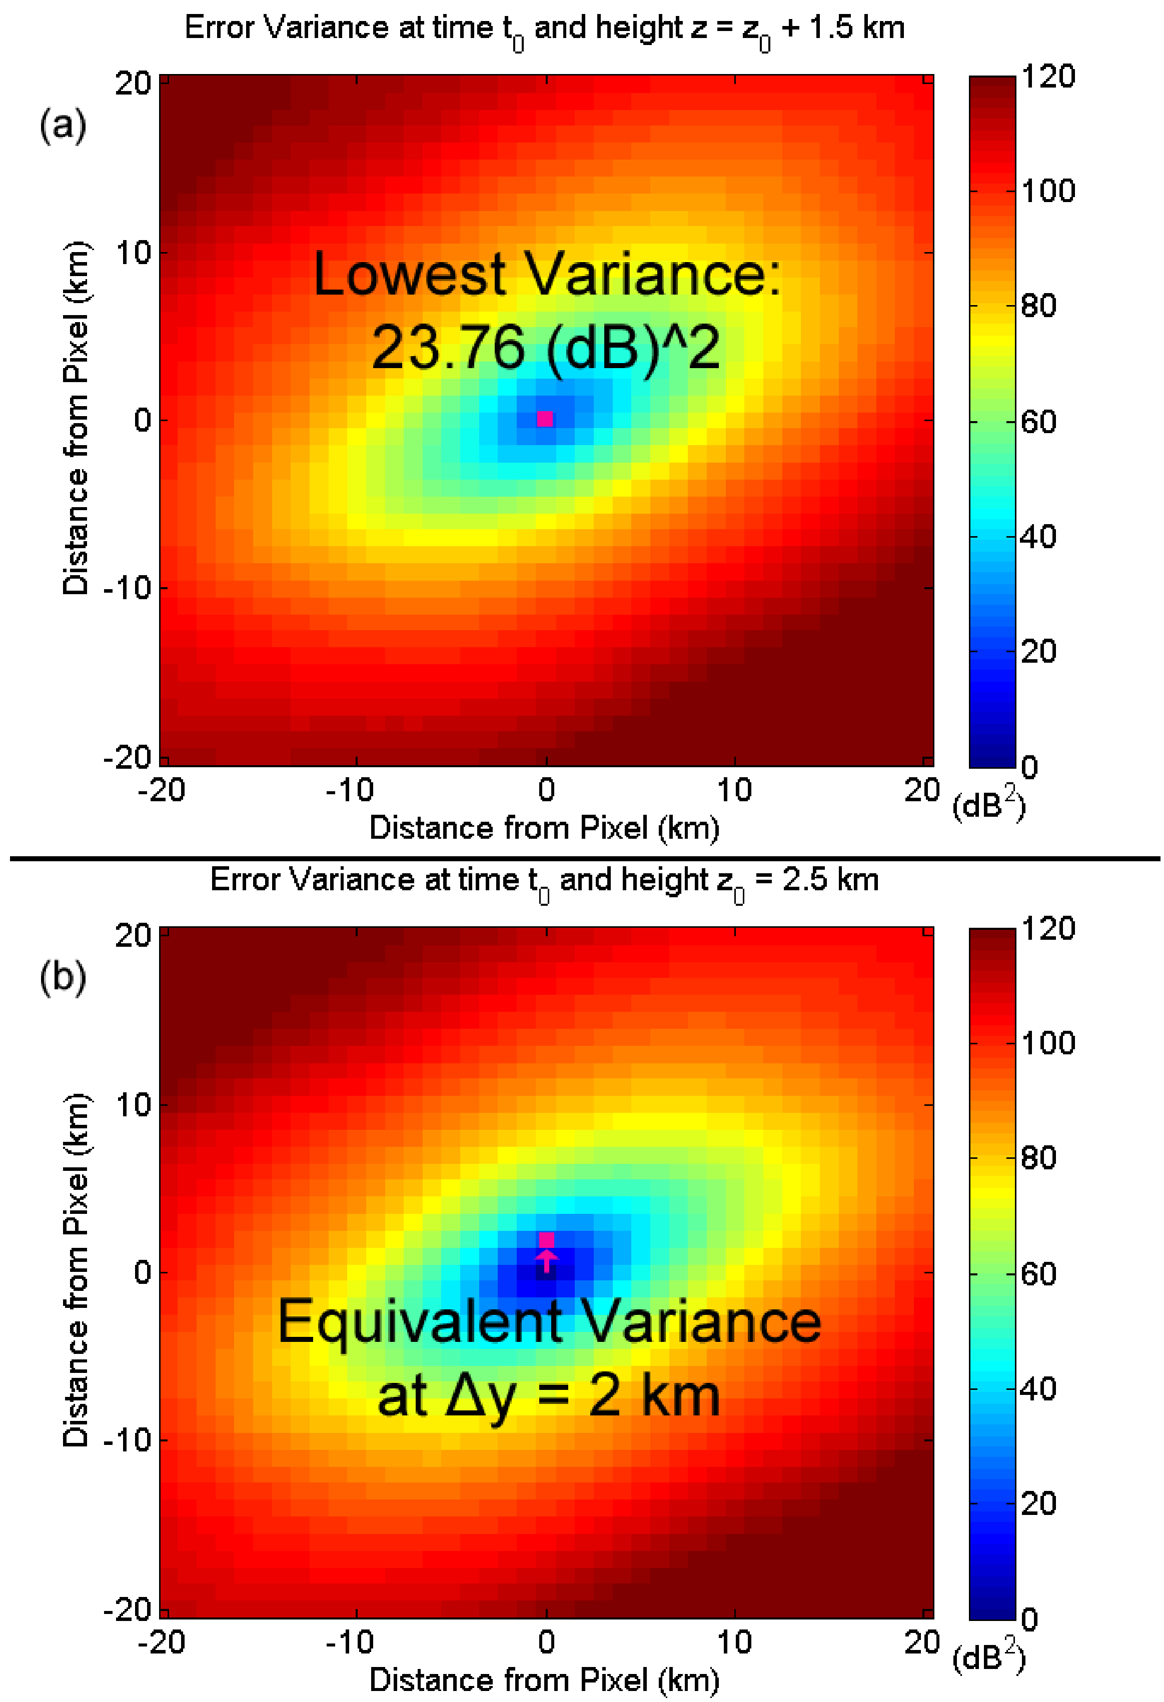
Error Variance at time t_0 and height $z = z_0 + 1.5$ km
(a)
Lowest Variance:
23.76 (dB)^2
Distance from Pixel (km)
Distance from Pixel (km)
(dB^2)
Error Variance at time t_0 and height $z_0 = 2.5$ km
(b)
Equivalent Variance
at $\Delta y = 2$ km
Distance from Pixel (km)
Distance from Pixel (km)
(dB^2)

Figure 3.9: Demonstration of the calculation of error (i.e., variance) equivalence in height. Our reference height is selected to be $z_0 = 2.5$ km, and we are attempting to establish the equivalent horizontal distance for $z = z_0 + \Delta z$, where $\Delta z = 1.5$ km. Figure 3.9a shows the variance at a vertical distance $\Delta z = 1.5$ km from our reference height z_0 and identifies the minimum value of variance at that height: 23.8 dB^2. We then search the variogram at our reference height z_0 (Fig. 3.9b) for the pixel with the nearest possible value to 23.8 dB^2, which turns out to be at a distance $\Delta y = 2$ km from our reference pixel. Thus, at a reference height of $z_0 = 2.5$ km for this convective event, moving $\Delta z = 1.5$ km higher in the vertical direction is approximately equivalent to moving $\Delta y = 2$ km in the horizontal direction.

This simple process, with a horizontal resolution of 1 km, will evidently not provide exact equivalences, but a general trend can be established by repeating this procedure for all nine heights relative to the reference height. A simple best-linear fit is made on a plot of height versus horizontal distance, and so a simple weighting equivalence is established—in order to replace a height level with its equivalent horizontal distance, we apply the transformation in the form of the equation of the best-fit line. This process is illustrated in Fig. 3.10 for Case 1, a convective event, and for Case 8, a stratiform event. For the sake of clarity, let us assume that we want to replace a contaminated pixel with a single pixel from either a different height or from a different horizontal distance. At what horizontal distance does it start becoming beneficial to look for pixels at different heights instead? Figure 3.10 shows that, for a convective event, a single-pixel replacement form a height of 1 km will only be an improvement if we are looking more than 1.4 km away horizontally. For the stratiform event, this number climbs to 5.4 km—this is quantitative confirmation of the qualitative aspects of the variograms presented earlier in this chapter, which showed that the stratiform event would benefit less from the replacement of pixels in the vertical than would the convective event. Thus, if we have a stratiform precipitation event and a pixel of ground clutter, and our two options for replacing this clutter are at 1 km vertically and at 3 km horizontally, the horizontal pixel is a better option.

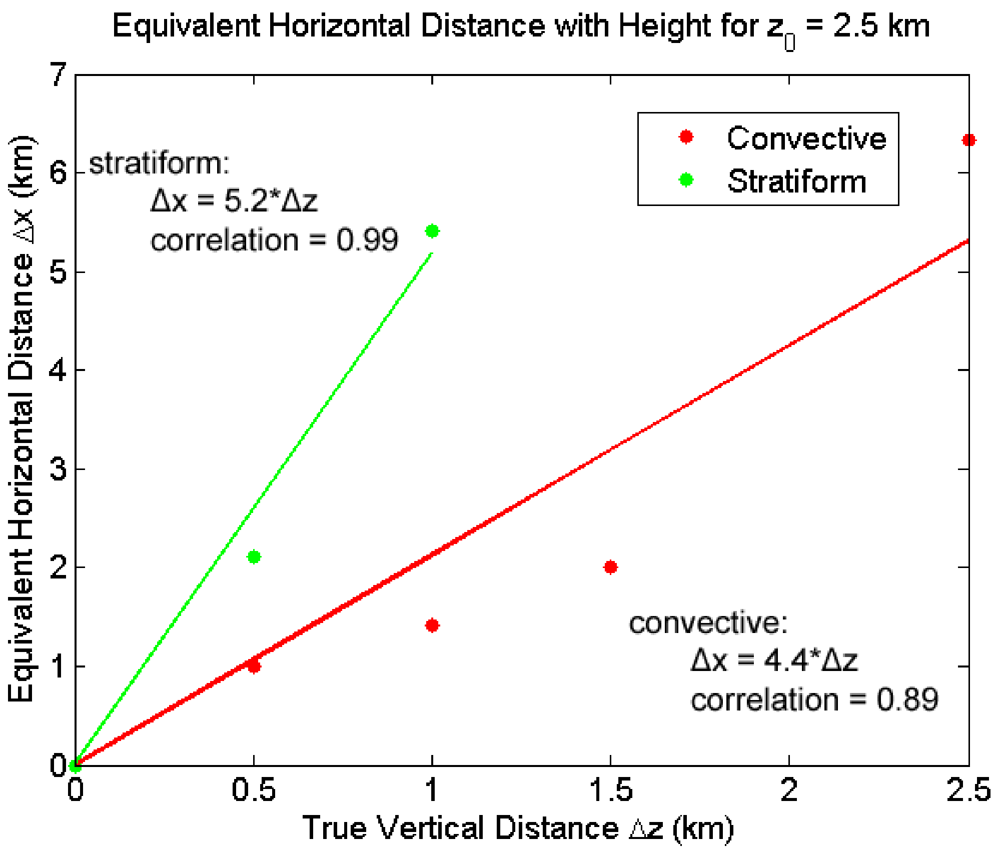

Figure 3.10: Convective (Case 1 only) and stratiform (Case 8 only) equivalent variance plot for height, starting from a reference height z_0 = 2.5 km. For instance, the equivalent horizontal distance corresponding to a true vertical distance Δz = 0.5 km is Δx = 1.8 km, and the value approximated by the trendline is Δx = 3.7*0.5 km = 1.85 km. The values for equivalent height corresponding to each true vertical distance are calculated as demonstrated in Fig. 3.9. The intercept is fixed to zero, since the equivalent horizontal distance calculated here is the absolute value of the horizontal distance, and so a negative value has no meaning. Likewise, the values of true vertical distance along the x-axis have been calculated as absolute values (for a vertical distance of 0.5 km, for

instance, an average was taken of the lowest variances for z_0 + 0.5 km and z_0 – 0.5 km).

To maximize the accuracy of this simple calculation, this process is repeated for each individual reference height. The same methodology is used to construct error equivalences in time, i.e., to determine at what horizontal distance it is actually preferable to take a replacement pixel from a different time instead. Refer to Appendix A for the tables of values and associated discussion.

Now that we have a solid understanding of the four-dimensional error structure for each of these nine events, as well as knowledge of error equivalences, we are capable of replacing a single bad pixel with its best possible (statistically speaking) neighbour. This is only a small part of the puzzle, however: the remainder of this study will be devoted to the development and evaluation of an algorithm that can replace large areas of ground clutter with optimal combinations of surrounding pixels.

Chapter 4: Towards an algorithm for information blending and correction of static ground clutter

4.1 Introduction and mechanics

The primary goal of this study is the creation of an automated algorithm that can replace known, stationary ground clutter in real-time using a combination of pixels that takes advantage of information at different spatial and temporal locations. An important requirement for this algorithm is that it should not smooth the data to the point where the extrema are no longer visible. In substantial rainfall events, these extrema are of paramount interest to forecasters, researchers, and the public alike, so it is essential that the final product avoid averaging data over too many pixels.

Four methods of increasing complexity are introduced in this chapter, ranging from the simplest two-dimensional nearest neighbour correction to an algorithm making use of ordinary kriging (OK) and the "smart" selection of a very limited number of pixels that can maximize sampling of independent spatial and temporal information and minimize smoothing. Finally, the methods are compared and statistics are calculated in order to evaluate each approach, balancing accuracy with the time necessary to regenerate products so as to incorporate near-future data.

In order to evaluate a given algorithm's ability to fill gaps due to ground clutter, we need a "reality" against which to compare that algorithm's final result. The simplest way of doing this is to create a region of false ground clutter based on reality, but translated so it does not intersect with the actual ground clutter: this way, we have access to the actual reflectivity field that this false ground clutter is concealing, and we are assured that the shape and size of this false ground clutter is realistic. The importance of doing this testing over a region of bad pixels rather than a single bad pixel is clear: in reality, the gaps in radar data that need filling are rarely confined to an isolated, 1 km × 1 km space. Figure 4.1 shows an area of false clutter that is based on the actual ground clutter pattern at McGill's radar (recall Fig. 2.2), transposed in such a way that the false clutter pixels never intersect with the real clutter pixels, thus assuring that there is always a reliable "truth" with which to compare the algorithms' results. The remainder of this chapter will demonstrate the results of replacing these false clutter pixels with different combinations of uncontaminated pixels.

Two specific scans will sometimes be used in the following sections as qualitative illustrations of the strengths and weaknesses of each gap-filling algorithm. The first of these reflectivity scans showcases an event in which some strong but localized convective cells have been obscured by the false clutter (Case 1; Fig. 4.2a-b), while the second is from an event in which the false clutter hides pockets of particularly strong stratiform precipitation (Case 8; Fig. 4.2c-d). More information on these particular samples will be given in the following sections. Note that, in many of the discussions that follow, focus will be given to the highly variable and horizontally inhomogeneous convective event as it is the more challenging of the two to simulate. Improvements in the algorithms' convective results are mirrored by improvements in the stratiform results, albeit less dramatically.

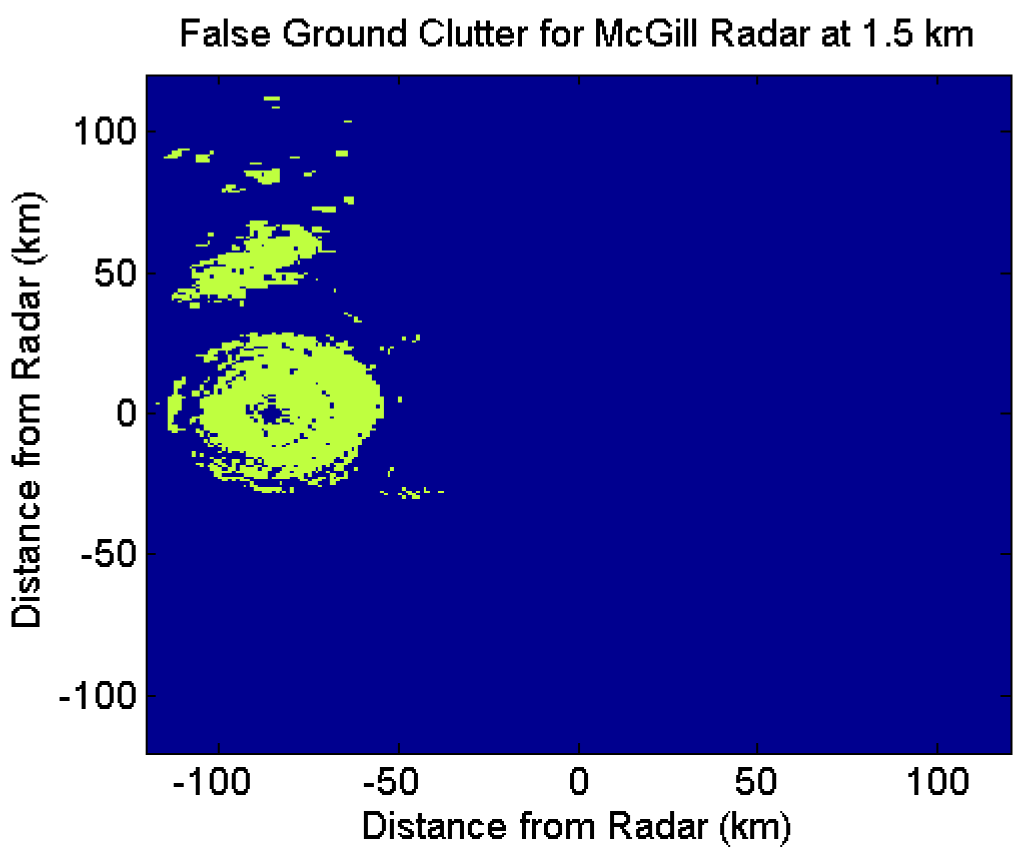

Figure 4.1: Depiction of the false ground clutter at the McGill radar for the 1.5 km CAPPI height (cf. Fig. 2.2, which is the actual ground clutter). The regions of green pixels, each one representing a 1 km × 1 km horizontal square, are

hereafter considered to contain "bad" data, and the purpose of the algorithms in the following sections will be to fill in these gaps with information that best reflects the true meteorological situation. The domain is 240 km × 240 km.

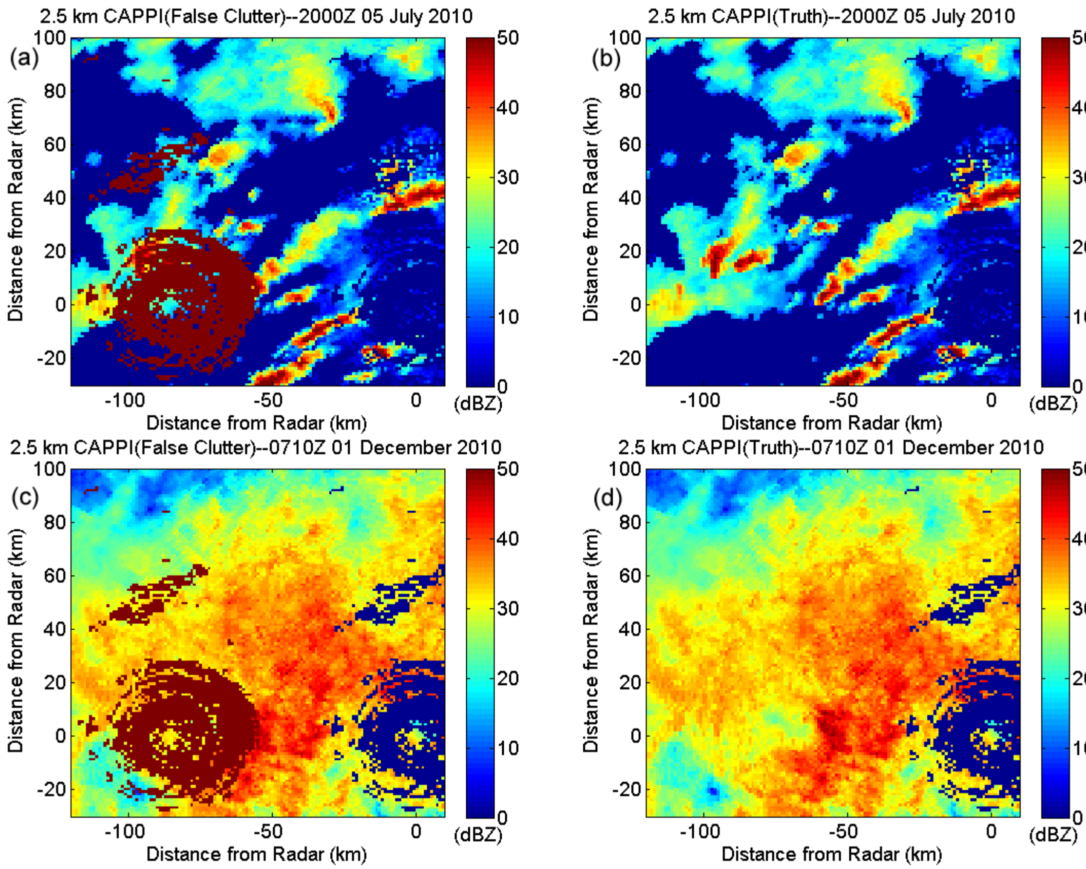

Figure 4.2: Sample radar scans to be used in qualitative algorithm evaluation for a convective case (a-b) and a stratiform case (c-d). The dark red area highlighted in the left column is the false clutter to be replaced by the following algorithms, while the right column shows the "truth" in this region, i.e., the actual recorded reflectivities. Note that these images are a subset of the 240 km × 240 km domain: the replaced pixels—the same as those shown in Fig. 4.1—are primarily

located about 80 km west of the radar (the blocked-out dark blue pixels in the lower right of each figure are the actual ground clutter centred around the radar).

4.2 Nearest neighbour algorithm

The first, and simplest, type of correction algorithm is a nearest-neighbour approach limited to two dimensions (hereafter NN2); that is, in seeking to replace our pixels of ground clutter, we use only the reflectivity data provided by a single reflectivity scan at a single vertical level. NN2 involves single-pixel replacement—rather than replacing our contaminated pixels with weighted combinations of pixels from their surroundings, NN2 will only make one-to-one replacements.

The process for NN2 is straightforward:

1 Identify all ground clutter pixels. In order to have a "truth" to compare with the algorithm output, for this exercise, we assign the algorithm the task of replacing each of the false ground clutter pixels depicted in Fig. 4.1. In reality, the NN2 would have as its input an array of archived pixels representing the radar's true ground clutter, as in Fig. 2.2.

2 For the first ground clutter pixel, $Z_0(x_0, y_0, z_0, t_0)$, search its surroundings for the nearest non-contaminated pixel, $Z_1(x, y, z_0, t_0)$, where $x = x_0 + \Delta x$ and $y = y_0 + \Delta y$, and "nearest" refers to the lowest available value of $r = \sqrt{[(\Delta x)^2 + (\Delta y)^2]}$.

3 Assign Z_0 the reflectivity value of the pixel Z_1.

4 If more than one pixel exists at the same horizontal distance from the pixel to be replaced, take the unweighted mean of their reflectivities.

5 Repeat steps 2-4 for each ground clutter pixel. Note that pixels that have already been replaced are not candidates for the replacement of other pixels: only true data will be candidates for nearest neighbour pixel replacement.

Figure 4.3 gives a graphical illustration of the performance of this two-dimensional nearest neighbour algorithm in (a) the convective case, and (b) the stratiform case. The dark lines passing through each plot show a "perfect" algorithm result, in which the reflectivity values suggested by the algorithm exactly match the reflectivity values present in reality. For the convective case, while there is an apparent grouping of these data points near the 1:1 line, there is also considerable spread, including values along the ordinate and abscissa where, for example, there is a pixel with a true reflectivity of 50 dBZ that the algorithm replaced with a 0 dBZ pixel, likely right at the edge of a convective cell. This relatively poor performance is to be expected of such a simple, horizontal-only approximation of a series of isolated convective storms with strong vertical development. This result should easily be beaten with the more complex algorithms to be described in the following sections.

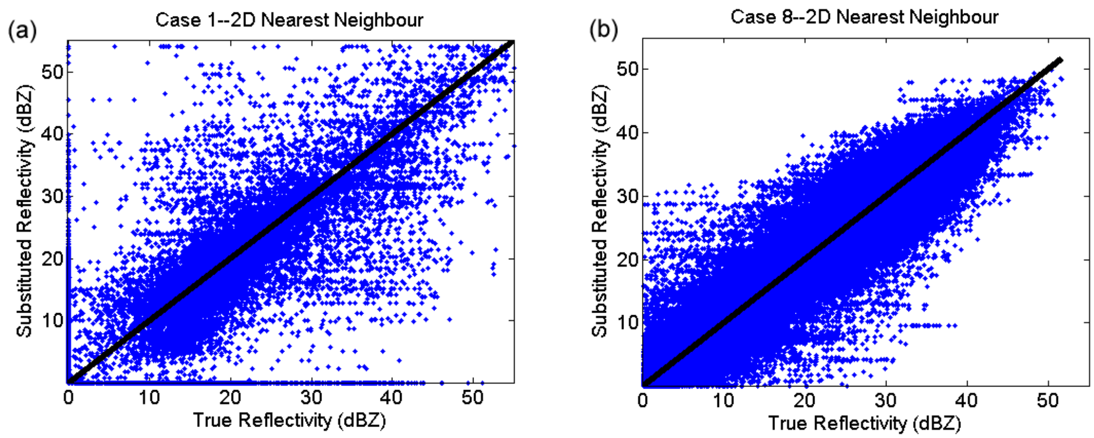

Figure 4.3: Scatterplots of true reflectivity values versus those generated by the NN2 algorithm ("substituted"). (a) Convective event (Case 1), (b) Stratiform event (Case 8). The black diagonal represents a 1:1 correlation between true reflectivity and substituted reflectivity—i.e., along this line the algorithm perfectly reproduces the actual reflectivity. Each point plotted corresponds with a single pixel of the false ground clutter depicted in Fig. 4.1.

Note that the scatterplot for Case 8 (Fig. 4.3b) shows a much more regular grouping around the "perfect" result, with fewer pixels directly along the x- or y-axes. The algorithm's improved performance reflects the stratiform event's stronger horizontal homogeneity: it is more likely that the nearest pixel in the horizontal will be an accurate substitution than it would be in our convective example. Note also that both plots show approximately the same number of pixels above and below the 1:1 line: there is no strong bias in the results, and so the nearest neighbour algorithm neither consistently overestimates nor consistently underestimates these reflectivity data. This is reflected by the fact that the mean reflectivity error for the convective event is only -0.0093 dB, and the mean error for the stratiform event is only 0.0112 dB.

To examine the results of this nearest neighbour algorithm in a more qualitative sense, two specific scans will be examined, one each from the convective and stratiform events (Fig. 4.4). These scans were selected to provide clear illustrations of the algorithm's performance in two different settings: the convective event (Fig. 4.4a-b) showcases the algorithm's ability to infer the presence and extent of relatively isolated convective cells, while the stratiform event (Fig. 4.4c-d) demonstrates its capacity for highlighting regions of enhanced or suppressed intensity within a relatively homogeneous precipitation field. These two scans will be used consistently throughout the more complex algorithms of future sections to provide a reference of the applied, physical effects of these algorithms on the derived reflectivity field.

Note that, in both the convective and stratiform precipitation events, the pixels replaced by the NN2 algorithm are readily apparent, especially in the region marked by the black circle in Fig. 4.4a, in the form of blocky areas of reflectivity. This is an expected result of this relatively limited algorithm for which there will often be no other choice but to use the same pixel to fill multiple gaps. The relative lack of bias is also apparent here: although the NN2 algorithm frequently misses the edges of precipitation, it can also infer much larger cells than occur in reality, if the nearest pixel to a particular gap has a relatively high reflectivity value (this is most apparent in Fig. 4.4a, about 50 km W of the radar).

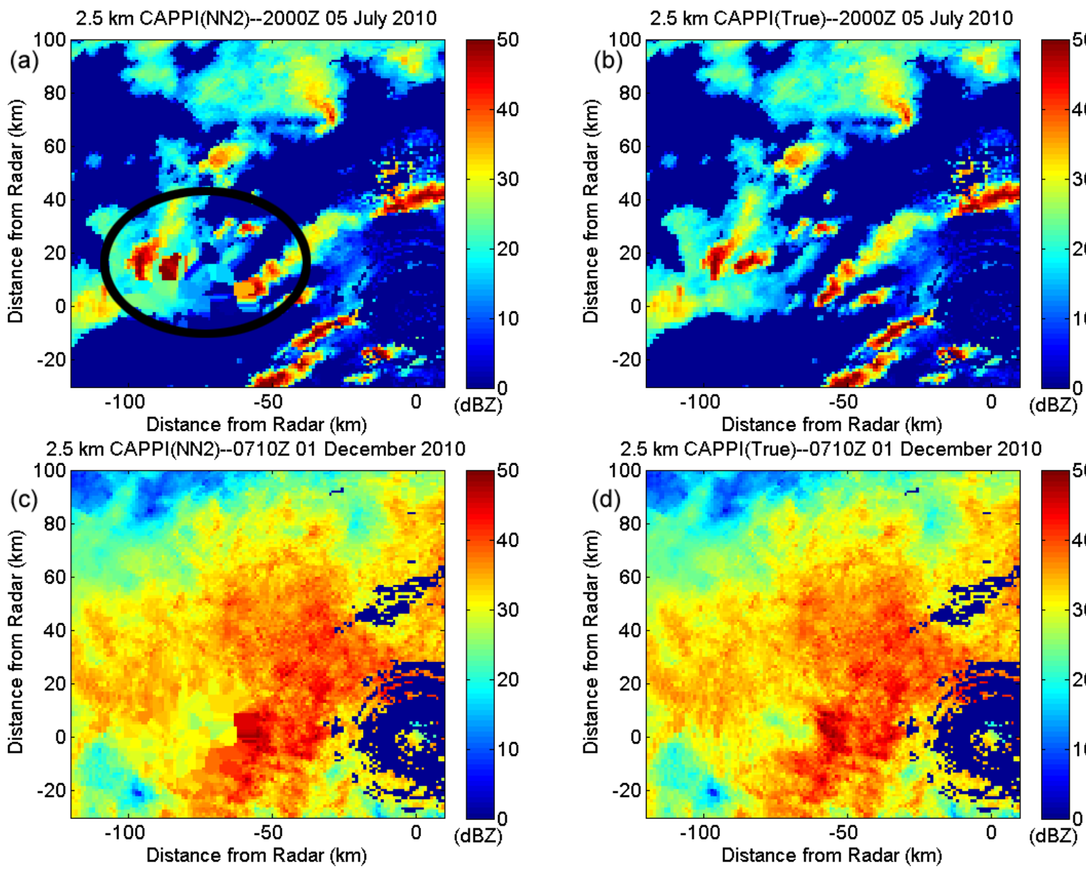

Figure 4.4: Examples of NN2 (2D Nearest Neighbour algorithm) performance for a convective event (a-b) and a stratiform event (c-d). Images in the right column are the "truth", i.e., the actual recorded reflectivities (see the right column of Fig. 4.2), while images in the left column show the reflectivity after the holes formed by the false ground clutter pixels have been filled by the NN2 method.

While this simple algorithm does provide results that are measurably better than *not* attempting to fill the gaps at all, the standard deviation from reality of these algorithm-derived reflectivity values is over 5 dB for the convective event and over 3 dB for the stratiform event. It should thus be fairly straightforward to formulate an algorithm to outperform this result and provide a clearer quantitative picture of the reality obscured by ground clutter. The next step in this attempt to create a more accurate algorithm is to bring in the variance data derived in the previous chapter.

4.3 Best pixel algorithms

The next level of complexity incorporates the variograms created in Chapter 3. The best pixel approach (hereafter BP) is again a single-pixel replacement algorithm—that is, it replaces a contaminated pixel with a single pixel from its surroundings—but this time, it makes that selection based on the statistics gathered. The procedure is as follows:

1 Identify all ground clutter pixels.

2 For the first ground clutter pixel, Z_0, examine its surroundings via the variograms created in Chapter 3, looking for the lowest variance that still corresponds to an uncontaminated pixel. Call this pixel Z_1.

3 Assign Z_0 the reflectivity value of Z_1.

4 Repeat steps 2-3 for every ground clutter pixel.

In its 2D form, we expect this approach to perform similarly to the 2D nearest neighbour method, but unlike the NN2 algorithm, we will allow the hunt for replacement pixels in the BP algorithm to be expanded to 3D (x,y,z), 3.5D (x,y,z, past times only), and 4D (x,y,z, all times). This expansion to higher dimensions should provide a substantial improvement over the limited NN2 algorithm.

Figure 4.5 clearly demonstrates the qualitative difference between the nearest neighbour and best pixel algorithms, especially in the convective example (top row). While the blockiness associated with single-pixel replacement is still unavoidably present, note that the results of the best pixel algorithm show stretching along the SW-NE axis (Fig. 4.5b), echoing the shape of the variogram for this event (recall Fig. 3.2). This effect is less immediately obvious for the more uniform stratiform precipitation.

Increasing the dimensionality of the best pixel algorithm essentially expands the possible pool of pixels from which to select the single replacement pixel, which reduces some of the observed blockiness. Incorporating the height dimension, for instance, reduces the standard deviation (of replacement pixels compared with reality in the convective event) from 5.4 dB to 5.2 dB. Incorporating pixels from different times when they happen to be the best available choice brings the standard deviation down considerably further, to 4.1 dB, although there is little difference between the 3.5D and 4D results: widening the search for replacement pixels to include near-future data has limited effect once past data have already been included. This simple algorithm can be thought of as a proof-of-concept for the value of considering data from different heights, and, especially, different times.

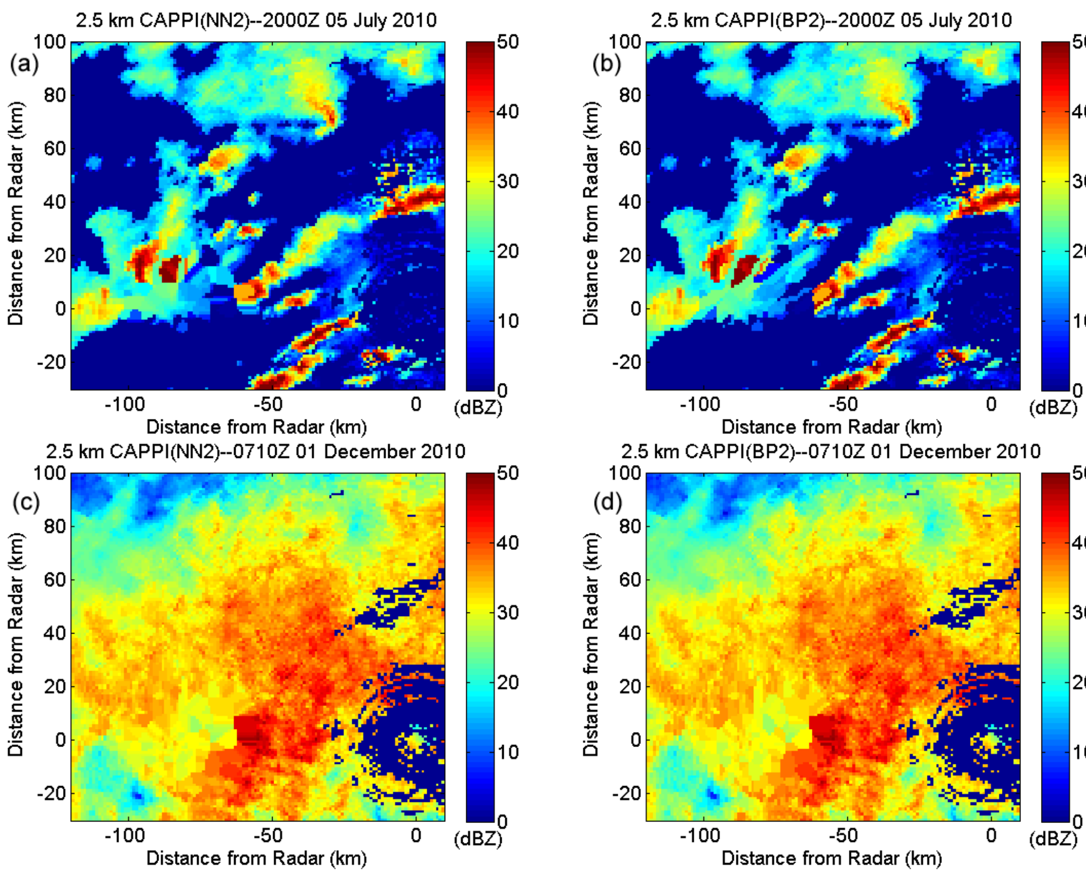

Figure 4.5: As in Fig. 4.4, but for a comparison of the NN2 (2D nearest neighbour algorithm; left column) and the BP2 (2D best pixel algorithm; right column). The top row is the convective event, while the bottom row is the stratiform event.

Both the nearest neighbour and best pixel algorithms use single-pixel replacement methods, that is, each pixel contaminated by ground clutter is replaced by a single pixel from elsewhere in the domain—the only difference is in the method used to select that pixel. As no other simple, single-pixel replacement strategies would be likely to make an appreciable difference to the final result, any improvements should emerge as a result of using *multiple*-pixel replacement instead, i.e., algorithms in which each pixel contaminated by ground clutter is replaced with some sort of weighted average of pixels from its surroundings rather than just another single pixel. The following two sections introduce the concept of ordinary kriging, and evaluate its use as a means of creating weighted averages to fill in the gaps caused by ground clutter.

4.4 Simple ordinary kriging algorithms

Ordinary kriging (OK), summarized in Appendix B, requires a number of pixels to be selected as possible replacement candidates, assigns each of those pixels a weight based on the mean error structure calculated in Chapter 3, and finally takes their weighted average. Attempting to average n pixels using ordinary kriging requires the inversion of an $(n+1) \times (n+1)$ matrix, which can easily become a computationally expensive process. Therefore, limiting the number of pixels to be kriged is an important challenge in developing a pixel-replacement algorithm that uses ordinary kriging. A lower number of pixels will also reduce the smoothing of extreme values that occurs with the averaging of many pixels (Section 4.6). In this and the following ordinary kriging methods, if N is the number of dimensions under consideration, the number of pixels to be kriged is selected to be 2N. This number was deemed to be an acceptable compromise between the lack of spatial information when there are too few pixels, and the over-smoothing and computational costs that arise with too many pixels. The problem has then been reduced to determining how best to select these 2N pixels, and two methods will be described in this and the following section. This section's simple OK approach automatically selects the 2N pixels with the lowest variance, while the smart OK approach (Section 4.5) attempts to select its 2N pixels in such a way that they provide a more complete sampling of the independent spatial and temporal information available.

The simplest method of incorporating ordinary kriging into the ground clutter correction process is as follows:

1 Identify all ground clutter pixels.
2 For the first ground clutter pixel, Z_0, use the best pixel algorithm described in Section 4.3 to select the 2N pixels with the lowest variance, where N is the number of dimensions the particular algorithm considers (2D, 3D, 3.5D, or 4D).

3 Using the variograms created in Chapter 3 as input, use ordinary kriging, as described in Appendix B, to build a weighted average Z of the reflectivities of our 2N pixels.

4 Replace Z_0 with the value Z.

5 Repeat for all ground clutter pixels.

The result is an algorithm that replaces ground clutter pixels with a weighted average of the lowest-variance pixels nearby, so we expect a certain amount of smoothing to occur in comparison with the best pixel algorithm described in the previous section. Figure 4.6 shows the dramatic qualitative improvement that occurs when multiple-pixel replacement is used in lieu of single-pixel replacement. Even looking only at the two-dimensional version of the simple ordinary kriging algorithm (Fig. 4.6b), when compared with the two-dimensional best pixel algorithm (Fig. 4.6a), it provides a much less blocky and more qualitatively realistic depiction of the true meteorological situation being obscured by the regions of ground clutter. The stretching along the SW-NE axis of the system is still apparent in the simple kriging algorithm's results, since it is, after all, still reliant on the variograms developed in Chapter 3 and will thus reflect their shape.

The higher-dimensional versions of the simple ordinary kriging algorithm provide added value. For the sake of comparison, two convective cells have been highlighted in Fig. 4.6a: the cell to which the white arrow points will hereafter be referred to as Cell 1, and the cell to which the red arrow points will hereafter be referred to as Cell 2. First, while the two-dimensional simple ordinary kriging algorithm's output is less blocky than that from the best pixel algorithm, the general shape and magnitude of both Cell 1 and Cell 2 in SiOK2 (Fig. 4.6b) are rather similar to BP2's result (Fig. 4.6a). Looking at the simple ordinary kriging algorithm when the height dimension is brought into play (Fig. 4.6c), the overdone extent of Cell 2 has been scaled back to something more closely approximating the true shape and size of that particular cell (cf. Fig. 4.4b). The addition of the time dimension in Fig. 4.6d stretches Cell 1 in the correct direction, very closely approximating the true shape of that cell as well.

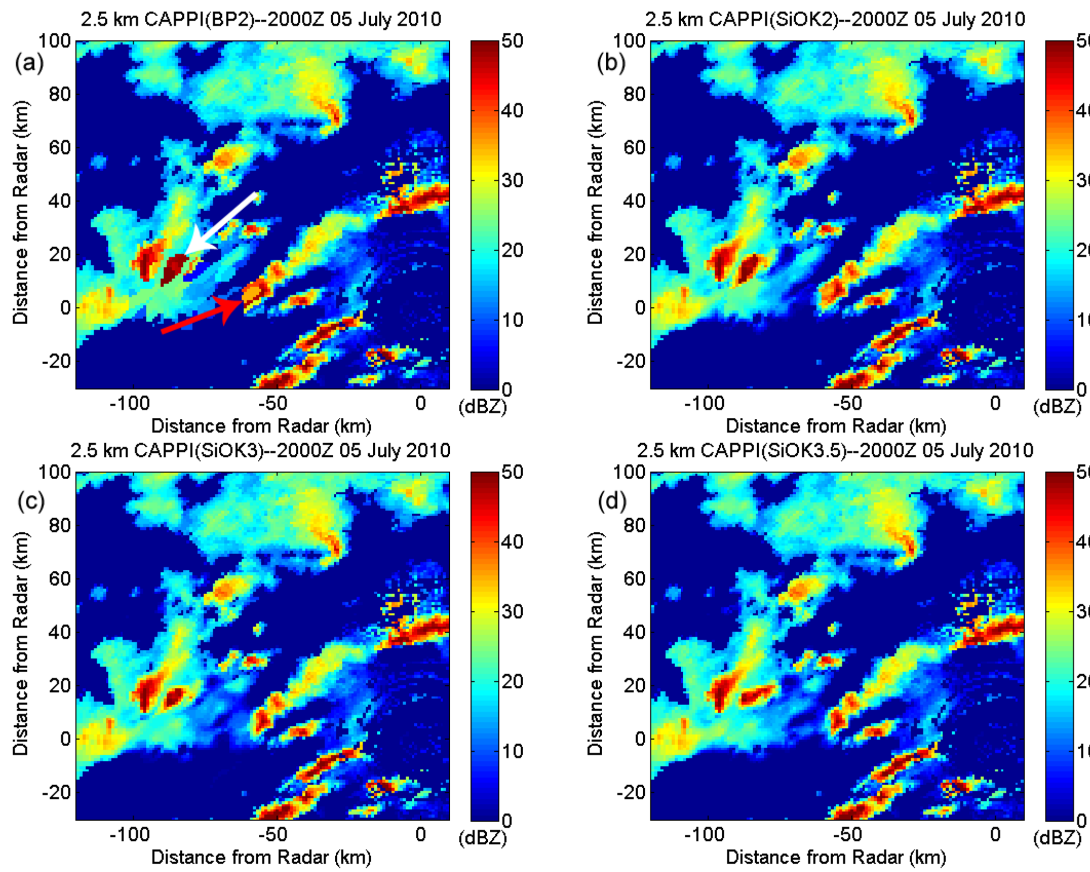

Figure 4.6: Comparison for the convective event to illustrate the benefits of the multiple-pixel selection approach. (a) Two-dimensional best pixel algorithm, with white arrow indicating cell 1 and red arrow indicating cell 2. (b) Two-dimensional simple ordinary kriging algorithm. (c) Three-dimensional simple ordinary kriging algorithm. (d) Three-and-a-half-dimensional (past times only) simple ordinary kriging algorithm.

The quantitative differences between these algorithms' results and the true reflectivities concealed by the false ground clutter bear out these qualitative improvements: going from a two-dimensional best pixel algorithm to a two-dimensional simple ordinary kriging algorithm, the standard deviation from the true values diminishes from 5.4 dB to 4.4 dB. The advantage of a multiple-pixel replacement scheme is apparent: the magnitude of this decrease in standard deviation is similar to that obtained when switching from a two-dimensional single-pixel replacement scheme to a fully four-dimensional one, at a much lower computational cost. The inclusion of pixels from different heights and times into the simple ordinary kriging algorithm improves the results still further, down to a 4.1 dB standard deviation.

In spite of these encouraging results, the simple ordinary kriging algorithm is by necessity making use of redundant information in its weighted averaging: it could well grab the values from two immediately adjacent pixels for kriging, with no particular net gain of information. In the following section, the final algorithm of this study will be developed: a "smart" ordinary kriging algorithm that prevents the selection of redundant information and encourages a sampling of data that represents the overall structure of the precipitation system.

4.5 Smart ordinary kriging algorithms

As discussed in the previous section, the simple ordinary kriging approach replaces pixels based on the mean variance structure derived in Chapter 3. This simple kriging method, however, merely considers candidate pixels as an ordered list, sorted from lowest to highest variance, without taking into account that the top pixels on that list may all provide redundant data. For example, given the continuous nature of radar reflectivity, if the best two available pixels are right next to each other, they are very likely to have nearly the same value, and the simple ordinary kriging algorithm will select these two pixels rather than taking just one and seeking out a different pixel that is more likely to provide independent information. The "smart" ordinary kriging approach described here is an attempt to mitigate these pixel-selection issues as much as possible by disallowing the selection of adjacent or nearby pixels, all without increasing the number of pixels selected (and hence the computational cost for 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, then, to 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 centred 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.

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.

We begin, as we did in the simple ordinary kriging algorithm, by selecting the best uncontaminated pixel using the variograms constructed in Chapter 3. Once this pixel (Z_1) has been selected, the axis linking it to the contaminated pixel Z_0 to be replaced is named as our first axis. We look in the negative direction along this axis and select a second pixel on the opposite side of Z_0, doing a search within a 30-degree arc for the best possible uncontaminated pixel (as determined by the variograms created in Chapter 3). This region is marked by the green lines shown in Fig. 4.7. Next, we select a second axis orthogonal to the first, and perform a similar check within a 30-degree angle in order to determine the next two pixels (see the blue lines in Fig. 4.7). This process is repeated until we have 2N pixels selected, where N is the number of dimensions under consideration—in this case, the number of orthogonal axes considered. The choice of using a 30-degree arc (as opposed to a 45-degree arc or a 15-degree arc) was somewhat arbitrarily deemed the most acceptable tradeoff between over-limiting a search region and opening the search region so wide that redundant pixels could still be selected. For instance, two 45-degree pixel search arcs intersect along a line, so it would be possible for two pixels to be selected along this line, providing redundant information—which is precisely what the smart ordinary kriging method hopes to avoid.

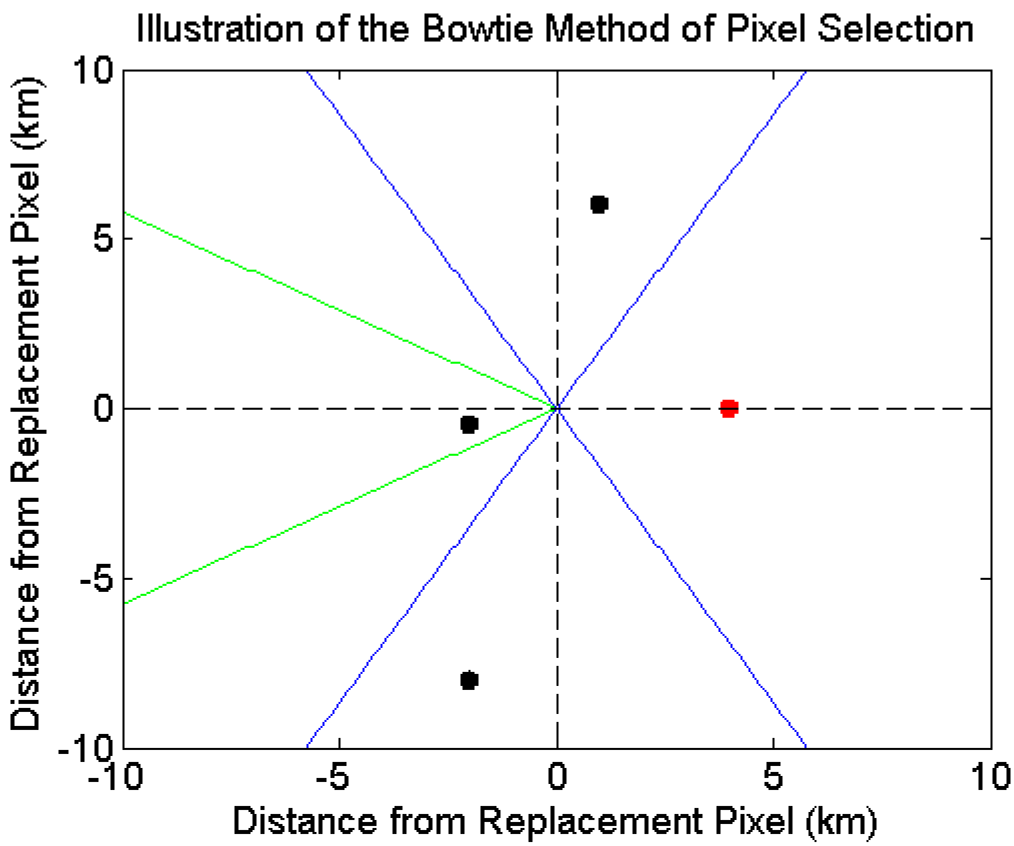

Figure 4.7: Illustration of the "bowtie" pixel selection method for the 2D case. The pixel with the lowest available variance (red point) is the first of the four pixels selected, and a system of axes (dashed lines) is built around it, based on the position of this best pixel with respect to the reference pixel to be replaced. 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. This process is trivially

generalizable to higher dimensions with the conversion of height and time to equivalent horizontal distances using Tables A.1 and A.2 (Appendix A).

Note that we will have to calculate these 30-degree arcs on unusual planes that include time and height as well as horizontal distance. Recalling the error equivalences calculated in Tables A.1 and A.2, by using these tables to relate the error structures of height, time, and horizontal distance, we are able to convert heights and times to horizontal distances, and then use these horizontal distances to calculate 30-degree angles as in Fig. 4.7.

The general process for the smart ordinary kriging approach is thus as follows:

1 Identify all ground clutter pixels.
2 For the first ground clutter pixel, Z_0, use the variograms created in Chapter 3 to determine the single uncluttered pixel Z_1 with the lowest variance (as in the best pixel algorithm described in Section 4.3).
3 Draw a 30-degree arc extending back from an axis A_1 passing through Z_1 and Z_0. Using the variograms, find the uncluttered pixel Z_2 with the lowest variance that lies within this 30-degree arc.
4 Including the initial axis A_1, find N orthogonal axes, where N is the number of dimensions under consideration. Again, draw 30-degree arcs in the positive

and negative directions along these axes and find the lowest-variance uncluttered pixels that lie within those 30-degree arcs.

5 While calculating 30-degree arcs, where necessary, convert heights and times to equivalent horizontal distances using the numbers derived in Appendix A.

6 If no uncluttered pixels are available within a given arc, do not select a pixel for that dimension. Thus, the total number of pixels selected will be somewhere between 1 and 2N, inclusive.

7 Using the variograms created in Chapter 3 as input, use ordinary kriging to build a weighted average Z of the reflectivities of our selected pixels.

8 Replace Z_0 with the value Z.

9 Repeat for all pixels contaminated by ground clutter.

The final result will be that each contaminated pixel has been replaced with a weighted average of surrounding pixels, which have been selected to include the most non-redundant spatial information.

Does this method provide an improvement over the simple ordinary kriging approach of the previous section? Even when the dataset is limited to only two dimensions, the answer is yes: the simple version of the 2D ordinary kriging algorithm (without the bowtie pixel selection method) has a 4.44 dB standard deviation, while the "smart" version has 4.40 dB. In fact, the smart ordinary kriging method has a slightly lower standard deviation than the simple ordinary kriging method in 2D, 3D, 3.5D, and 4D. Does that suggest that it is unambiguously the best choice of algorithm to use for correcting ground clutter? Should this smart pixel-selection method, based on accurately sampling time-space variability, always supplant the simpler approach of blindly picking the lowest-variance pixels? As always, the data user must weigh the potential benefits against the added complexity of implementing the smart kriging method's more elaborate pixel selection procedure. These benefits are quite small: as described above, the smart ordinary kriging algorithm provides only a 1% lower standard deviation when compared with simple ordinary kriging at 2D (4.44 dB versus 4.40 dB), and this improvement is not substantially larger even in the fully 4D algorithms.

If time is short (as in many operational situations) or computational power is limited, the added bulk of the more complicated pixel selection process is not ideal, given such a paltry reward. A less computationally expensive method of overcoming the simple ordinary kriging method's tendency to pick redundant pixels is simply to select *more* pixels: upping the number of pixels kriged in the 2D simple ordinary kriging method from 4 to 8 reduces its standard deviation from 4.44 dB to 4.26 dB, making it a 3% *improvement* over the 2D smart ordinary kriging algorithm while still taking slightly less time to run. The disadvantage of this approach is that averaging over more pixels tends to smooth away extreme values, an issue that will be discussed in the following section. Given an ideal situation, however, with high processing power and/or few time constraints (e.g., regenerating earlier data to use in a research case study), the slight improvement of the smart ordinary kriging algorithm is an improvement all the same. For example, in the 2D case, as long as those first four pixels are selected using the bowtie pixel selection method of Fig. 4.7, the smart kriging algorithm's performance will always edge out the simple kriging algorithm's performance.

Figure 4.8 shows, in the form of scatterplots, the progress made from the two-dimensional nearest neighbour algorithm through to the fully four-dimensional smart ordinary kriging algorithm. As is reflected by the 1.8-fold reduction in standard deviation, there is less spread around the 1:1 line for the four-dimensional smart ordinary kriging algorithm than for the two-dimensional nearest neighbour algorithm. However, while the nearest neighbour algorithm showed no particular bias (see discussion in Section 4.2), in the four-dimensional smart ordinary kriging algorithm, a clear majority of the pixels lie under the 1:1 line (i.e., the substituted reflectivity values are often lower than the true reflectivity values), and there are relatively few pixels well above the 1:1 line, unlike with the NN2 algorithm. We would expect the more complex multiple-pixel algorithms to have a tendency to underestimate reflectivity values: this is the smoothing effect of averaging, blurring out the extrema, and this effect would have been even more pronounced if we had used more than 2N pixels in the ordinary kriging algorithm.

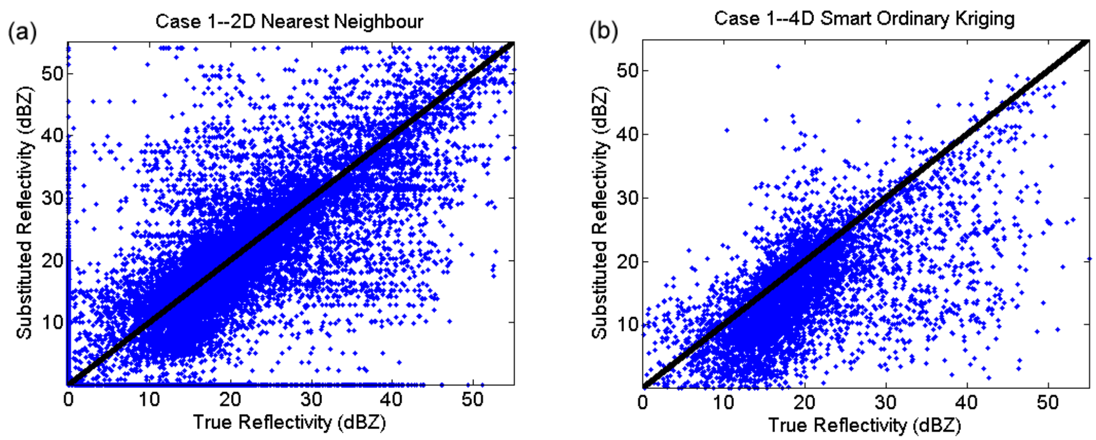

Figure 4.8: As in Fig. 4.3, but for (a) the two-dimensional nearest neighbour algorithm, and (b) the four-dimensional smart ordinary kriging algorithm. Both scatterplots are for the convective case. The stratiform case (not pictured) did show more bunching around the 1:1 line for the 4D smart ordinary kriging algorithm, as well as a slight bias toward the bottom half of the plot, but the contrast was less visually striking due to the relatively strong performance of the 2D nearest neighbour algorithm in the stratiform case.

The following section evaluates the performance of these algorithms when it comes to creating more complex products derived from simple reflectivity imagery. Special focus will be placed on the top percentile of instantaneous rainfall rates alone, to evaluate how much of a detriment the aforementioned "smoothing" effect may be on the algorithms' performance when replacing the pixels with extreme values that signal particularly severe weather.

4.6 Algorithm performance in additional radar products

We have thus far limited the evaluation of the algorithms' performance to raw CAPPI reflectivity replacement, i.e., looking at how well they perform in real-time with the CAPPI radar scans forecasters would be observing as they are generated in an operational center. These raw data can also be used to create a variety of radar products related to rainfall rates and accumulations, and by evaluating the algorithms' performance with these additional challenges, the full range of potential applications for these algorithms can be appreciated. The following subsections evaluate the algorithms' performance for instantaneous rainfall rates and for one-hour rainfall totals. An additional test evaluates their performance for the top percentile of instantaneous rainfall rates only, a "worst-case" scenario associated with anomalously strong precipitation.

4.6.1 Rainfall rates and totals

The accurate conversion of radar reflectivity to rainfall rate is a topic that has engendered a great deal of discussion and debate in the research community (e.g., Uijlenhoet 2001, Marshall et al. 1955). For the sake of simplicity, and direct analogy with operational applications, the conversion formula used in this brief evaluation is chosen to be the same as that employed by the U.S. National Weather Service for their Weather Surveillance Radar 88 Doppler (WSR-88D) network, i.e., $Z = 300R^{1.4}$ (Fournier 1999), where Z is the reflectivity value and R is the rainfall rate in mm/hr.

The absolute differences between the algorithms' substituted instantaneous rainfall rates and reality show a similar trend to those in reflectivity, with slightly lower errors across the board for the stratiform event versus the convective event. Table 4.1 compares these results with the results of Sánchez-Diezma et al. (SD01; 2001) introduced in the first chapter. Recall that SD01's methodology involves using a simple thresholding approach to distinguish between convective and stratiform events, and then chooses its single-pixel replacement method accordingly. While the approaches introduced in this chapter outperform the SD01 method by a considerable margin (the smart ordinary kriging approach results in a sevenfold reduction in error), they have the added advantage of not having to distinguish between stratiform and convective events: the same algorithm is applied to both.

Table 4.1: Rainfall rate error comparison for Sánchez-Diezma et al. (SD01; 2001), the 2D nearest neighbour algorithm (NN2), and the 4D simple ordinary

kriging algorithm (SiOK4). See Section 1.3 and text for a more detailed description of the SD01 horizontal+vertical algorithm. NN2 and SiOK4 data are from the convective event. Stratiform event errors (not shown) are slightly lower than those for the convective event.

Algorithm Type	Mean Error in Rainfall Rate (mm/hr)
SD01 horizontal+vertical	1.04
NN2	0.71
SiOK4	0.15

These error values are extremely low—it is important to keep in mind that they are strongly skewed by the large regions without reflectivity, or with very low reflectivities, where even a 5 dB error in reflectivity corresponds only to a minute (<0.1 mm/hr) error in rainfall rate. Contrast this with the higher end of the reflectivity scale where, for instance, a 5 dB error could correspond to nearly 5 mm/hr of error in rainfall rate. The following subsection will address the issue of extreme precipitation values and evaluate the algorithms' performances in these specific situations.

As is the case in the reflectivity imagery of the previous section, the rainfall rates show a blockiness in the output of the single-pixel replacement algorithms (i.e., nearest neighbour and best pixel), and this blockiness is in turn smoothed out by the multiple-pixel replacement algorithms (i.e., simple and smart ordinary kriging). A relatively negative bias is apparent in both types of kriging, where the smoothing due to the algorithms' weighted averages results in an underestimation of extreme rainfall amounts. This bias is most obvious in the versions of the algorithms that take their averages over more pixels, so that the 2D versions of both the simple and smart ordinary kriging algorithms have mean negative errors of about 0.2 mm/hr, while their 4D analogues have mean negative errors closer to 0.4 mm/hr. There is thus an additional tradeoff inherent in increasing the number of pixels under consideration: while the inclusion of a larger number of pixels into the kriging process generally results in a lower *absolute* error, the *signed* error will show a stronger negative bias overall as extrema are smoothed out.

One-hour rainfall totals (calculated by determining the mean rainfall rate in mm/hr over a given one-hour period) should show a reduced error when compared with the instantaneous values. For instance, we would expect errors such as those associated with poor algorithm performance along the edge of a tight reflectivity gradient to be smoothed out over time as the precipitation moves out of the region of ground clutter. Table 4.2 illustrates that this is precisely what occurs: the errors for one-hour totals are generally about half of those obtained for instantaneous rainfall rates. In addition, the inclusion of data from the height and time dimensions, as well as the shift from single-pixel replacement algorithms to multiple-pixel replacement algorithms, both contribute to lower the absolute error in the one-hour rainfall total, just as they did for the instantaneous rainfall rates.

Table 4.2: Comparisons of absolute error for instantaneous rainfall rates and one-hour rainfall totals for the two-dimensional nearest neighbour algorithm. Values are provided for the convective event and the stratiform event.

Precipitation Type	**Instantaneous Rainfall Rate Error (mm/hr)**	**One-Hour Rainfall Total Error (mm)**
Convective	0.71	0.38
Stratiform	0.62	0.22

The rainfall rate data and the one-hour rainfall total data closely match the results obtained while looking at the raw reflectivity data, and while it is good to have confirmation of earlier results, no particularly novel information has been gained by shifting focus from reflectivity to rainfall rates, except perhaps to put these results into a more familiar physical context. For instance, in an hour's rainfall, an error of only 0.38 mm (as was observed for the simplest 2D nearest neighbour algorithm) during a convective event is vanishingly small for most practical purposes. However, given that this mean error has been calculated over a large region that contains quiescent and easy-to-predict areas that greatly outnumber the odd convective cell, it is important to make a closer inspection of the most extreme rainfall rates alone. To this end, in the following subsection, the algorithms' performance for the top percentile of rainfall rates is evaluated.

4.6.2 Top percentiles

While it has been established in previous sections that the multiple-pixel replacement algorithms (i.e., the simple and smart ordinary kriging algorithms) tend to outperform the single-pixel replacement algorithms (i.e., the nearest neighbour and best pixel algorithms) in reflectivity and in rainfall rate/rainfall total calculations, it seems likely that the smoothing associated with the multiple-pixel averaging may erase the contribution of the most extreme pixels. It is important to note that these extreme events are of particular interest to forecasters and researchers: after all, any improvement provided by the more complex algorithms is nearly meaningless for practical purposes if the pixels being replaced are only on the order of less than 10 dBZ in the first place. In order to investigate the possibility that these extreme values are being more strongly underestimated by the more complex algorithms, only the top percentile of rainfall rates will now be included when calculating the mean errors. For the convective case, this translates to evaluating algorithm performance only at rainfall rates greater than about 18 mm/hr, whereas for the stratiform case the top percentile begins at closer to 12 mm/hr.

Since we are considering only the top 1% of rainfall rates, we expect all algorithms to underestimate the results to at least some degree. We also expect that any algorithms that involve averaging over larger numbers of pixels will tend to smooth out these extreme values to a greater degree than the single-pixel replacement algorithms; that is, the negative bias will be most pronounced in multiple-pixel replacement algorithms such as the simple and smart ordinary kriging algorithms. Since modifying these kriging algorithms to incorporate data from different height levels and times also increases the number of pixels being averaged, the higher-dimensional versions of these algorithms should show a stronger negative bias than the 2D versions.

Table 4.3 displays an illustrative sample of the algorithms' performance for these extreme values of instantaneous rainfall rate. The "mean error" column refers to the mean *signed* difference between the algorithm's predicted value and the true value; in this column, a negative value indicates that the algorithm is consistently underestimating the true rainfall rate. The "standard deviation" column can be thought of as a measure of the *absolute* difference between a typical algorithm-generated value and the true value, so that lower values in this column correspond with more accurate predictions overall. Note that the negative bias is stronger for the smart ordinary kriging (multiple-pixel replacement) algorithms than for the best pixel (single-pixel replacement) algorithms; that is, the more complex algorithms actually have a harder time capturing the more extreme rainfall rates. Note, too, that the standard deviation for the smart ordinary kriging algorithm is higher than that for the best pixel algorithm, for the first time in this study: when considering these extreme reflectivity rates in isolation, there is no apparent advantage to choosing a more complex algorithm. The stratiform event (not shown) displays similar results, albeit at a lower magnitude due to its less extreme values of rainfall rate.

Table 4.3: Error comparisons for the convective case's top percentile of instantaneous rainfall rates. Results compared are for the two-dimensional best pixel algorithm (BP2), the two-dimensional smart ordinary kriging algorithm (SmOK2), the three-dimensional best pixel algorithm (BP3), and the three-dimensional smart ordinary kriging algorithm (SmOK3).

Algorithm	Mean Error (mm/hr)	Standard Deviation (mm/hr)
BP2	-11.61	34.99
SmOK2	-25.05	37.22
BP3	-13.18	35.96

SmOK3	-31.60	43.40

Does this result then suggest that, during extreme precipitation events, it would be preferable to set aside the more complex multi-dimensional ordinary kriging algorithms and return to basic 2D best pixel or nearest neighbour approaches? First, it is important to note from Table 4.3 that, while the best pixel algorithms have a much less substantial bias overall, their standard deviation is still extremely high, i.e., nearly at the same level as the smart ordinary kriging algorithms. This small minority of extreme rainfall rates may see a small relative drop in standard deviation (from 43 mm/hr to 36 mm/hr) when using the best pixel or nearest neighbour algorithms instead of the kriging algorithms, but that dubious improvement is not exactly a ringing endorsement for throwing aside the stronger *overall* performance of the more complex algorithms, as has been demonstrated throughout this chapter.

In addition, when the range of values to be evaluated is increased from the top 1% of rainfall rates to the top 5% of rainfall rates, a slight improvement is apparent: the standard deviations of the more complex kriging algorithms become marginally lower than those for the best pixel algorithms. That is, despite the persistent negative bias, the kriging methods still produce a better overall result than the simpler single-pixel methods, even for these extreme rainfall rates.

Earlier sections have shown that there are powerful quantitative benefits to making use of data from different times and heights, as well as using multiple-pixel combinations rather than focusing on single-pixel replacement. Looking at the specific problem of filling gaps that happen to contain extremely high rainfall rates, however, the added complexity of the higher-dimensional kriging algorithms tends to smooth away the highest values. In spite of this caveat, the fact remains that just because the simpler algorithms perform better than the kriging algorithms for extreme values does not mean they perform *well*. For practical applications, the slight improvement offered for these extreme values would not be likely to outweigh the better overall performance offered by the more complex kriging algorithms.

4.7 Summary

With information from reflectivity data, instantaneous rainfall rates, one-hour rainfall totals, and the top percentile of rainfall rates now available for each of these algorithms, the advantages and disadvantages of each algorithm can be summarized. Table 4.4 collects the points discussed over the course of this chapter for each type of algorithm, and also generalizes the effects of incorporating past data, near-future data, and data from different heights into these algorithms.

Table 4.4: Summary of algorithm advantages and disadvantages as discussed throughout Chapter 4. The final three rows generalize the effects of increasing the dimensionality of any of these algorithms. For the purposes of this table, "high" and "low" are relative terms referring to the other algorithm types or, for the last three rows, lower-dimensional algorithms. "Standard deviation" here refers to the typical difference between an algorithm-generated value of reflectivity or rainfall rate and the true value it is replacing; i.e., a low value indicates more accurate performance overall. "Typical error structure" refers to the variograms created in Chapter 3.

Algorithm	**Advantages**	**Disadvantages**
Nearest Neighbour	• low computational cost • low bias in extreme rainfall events	• highest standard deviations overall
Best Pixel	• low computational cost • low bias in extreme	• high standard deviations overall • requires knowledge of

	rainfall events	typical error structure
Simple Ordinary Kriging	• low computational cost • low standard deviations overall	• high bias in extreme rainfall events • requires knowledge of typical error structure
Smart Ordinary Kriging	• lowest standard deviations overall (marginally)	• high computational cost • high bias in extreme rainfall events • requires knowledge of typical error structure
Incorporation of Height Data (3D)	• reduces standard deviations overall (marginally) compared with 2D	• high computational cost
Incorporation of Past Data (3.5D)	• reduces standard deviations overall (dramatically) compared with 3D	• high computational cost
Incorporation of Near-	• reduces standard deviations overall	• high computational cost • requires wait-time and

Future Data (4D)	(marginally) compared with 3.5D	regeneration of product

Determining which algorithm is the most appropriate for use in gap-filling relies at least in part on the nature of the application in question: is there time enough to wait for near-future data to arrive in order to take advantage of the slight improvement a 4D algorithm provides over a 3.5D one? Are computational resources limited to the point where the smart ordinary kriging algorithm's minimal improvement over the simple ordinary kriging algorithm cannot justify its higher computational cost? Recall that in Section 4.5 it was shown that increasing the number of pixels selected in the simple ordinary kriging algorithm would result in a net improvement over the smart ordinary kriging algorithm without a dramatic increase in computation time. Is this modified version of the simple ordinary kriging algorithm a better option, given that averaging over more pixels will tend to smooth out extreme values still further?

Although the 4D smart ordinary kriging algorithm fills gaps most accurately overall, it is plagued by problems of diminishing return, i.e., past a certain point, costly changes to the algorithm are reflected by only very slight improvements in the results. For many applications with more limited resources in terms of both time and computing, the 3.5D simple ordinary kriging algorithm may represent the best of both worlds: it has the substantial advantage of incorporating data from different times without having to wait for new information to come in, it performs nearly as well as the smart version (see discussion in Section 4.5), and its computational cost is low.

Chapter 5: Summary and conclusions

Gaps in radar data are commonplace, whether they are caused by beam blockage due to obstacles, by attenuation due to areas of particularly heavy precipitation, or by persistent ground clutter. These holes present a two-part problem: first, contaminated data must be identified as such, which is an issue left to the sizeable sum of pre-existing research on the topic. Second, once an area of contaminated data has been identified, these gaps in the reliable data record must be filled in some way; the process of gap-filling of previously identified areas of ground clutter has been the driving topic of this study.

Using Montreal's J.S. Marshall Radar Observatory as a testing ground, a clutter mask was constructed to identify and blot out regions of ground clutter, using the long record of radar data for that facility. The data uncontaminated by ground clutter were then used to create variograms, which are statistical visualizations of the typical error that would result if any given pixel were replaced by a single pixel from its surroundings. These variograms provided vital information about the error structure of a variety of precipitation events, both convective and stratiform, at various stages of development. One major convective event (05 July 2005) and one major stratiform event (01 December 2010) were selected to showcase various steps along the way toward developing a real-time, operational algorithm that would be able to fill in the holes due to ground clutter (or, for that matter, any other known gap in the data). A region of realistic-looking false clutter was created and overlaid on the radar data, so that the algorithms' results could be compared with reality.

The gap-filling algorithm selected as a starting point was a two-dimensional setup that simply replaced any pixel contaminated by ground clutter with the nearest uncontaminated pixel. From there, two parallel (and complimentary) lines of improvement were pursued. The first focused on improving the fundamental methodology of the algorithm, first by selecting the replacement pixel using a variance criterion rather than mere proximity (the "best pixel" algorithm), then by incorporating a geostatistical averaging method, ordinary kriging, to combine either a small number of lowest-variance pixels (the "simple ordinary kriging" algorithm), or a selection of low-variance pixels chosen so as to minimize redundant information (the "smart ordinary kriging" algorithm). In test-runs of reflectivity data, for both the convective and the stratiform events, the more complex multiple-pixel replacement algorithms consistently outperformed the simpler single-pixel replacement algorithms. While the smart kriging was an improvement over the simple kriging, the difference between their performance was less than the major leap from the best-pixel approach to the simple kriging approach.

The second line of improvement focused on increasing the pool of data from which the replacement pixels for these algorithms could be drawn, by incorporating first the data from different CAPPI height levels (3D), then data from past radar scans (3.5D), and finally data from both past radar scans and future scans up to 25 minutes later (4D). The obvious disadvantage of the fully four-dimensional algorithms is that they require the user to wait the half-hour necessary for those near-future scans to come through: while regeneration of radar products is not a problem in some contexts, in real-time nowcasting applications it is less than ideal. However, there was shown to be little difference in skill between 3.5D and 4D algorithms, so that applications that must forego regeneration of their radar products would not see a major decrease in performance. A small jump in skill occurred between the 2D and the 3D algorithm (particularly for the convective test event, with its strong vertical homogeneity), but a much larger leap in skill occurred in both the convective and stratiform events when the time dimension was incorporated, to the point where reflectivity errors were nearly halved when compared with the 2D versions of the algorithm. Even when only past radar data is available, it is thus an extremely valuable source of information, and many radar products would be well-served by looking to earlier data for clutter correction.

The problem of seeking the best gap-filling strategy for CAPPIs seems to indicate either the 3.5D or (if near-real time regeneration of products is possible) the fully 4D smart ordinary kriging algorithm, with the caveat that the algorithm's kriging has a tendency to smooth out extreme values. This smoothing was found to most strongly affect only the top ~5% of rainfall rates, although it should be noted that all events examined in this study were substantial rainfall events; on more quiescent days, the strong negative bias of the top-percentile data will have less dramatic results. For instance, on a quieter day the maximum rainfall rate observed may be only 1 mm/hr, in which case an underestimation will make far less of a difference to most applications than if that maximum rainfall rate were, say, 20 mm/hr during a more substantial event.

The choice of the best algorithm to use for gap-filling depends in part on the limitations of the application in question. With limited time and computational resources, the 3.5D simple ordinary kriging algorithm will produce accurate results for nearly all purposes. If time is not pressing, radar products can be regenerated, enabling the use of near-future data as well as past data, which provides a slight improvement over the results with the use of past data alone. Finally, with more computational resources available, switching to the smart ordinary kriging algorithm will provide a slight (~2%) reduction in error.

The takeaway messages here are threefold: first, the ordinary kriging technique is a simple and versatile method that has seen very little use in radar meteorology thus far, and has mainly been limited to interpolation using broadly spaced datasets, 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 use this powerful method. Through the use of the bowtie pixel-selection method, however, this problem is rendered moot, by selecting a small number of representative pixels that can then be kriged for relatively little computational cost, broadening the range of applications for this method still further.

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.

Appendix A: Height and time error equivalence tables

The equivalences for each reference height, along with their R-squared values (nearer to 0 represents a poor correlation, while nearer to 1 represents a strong correlation), are shown in Table A.1. Calculations are done for each reference height to increase the accuracy of the final weighting equivalences.

Table A.1: Error equivalence coefficients for height, where Δz (in km) is the vertical distance of a given pixel from the reference pixel we wish to replace. These values can be used to calculate the horizontal distances Δx that are approximately equivalent to the vertical distances Δz in terms of variance. Values are calculated separately for each of the ten possible reference heights. The convective numbers are averaged over all six convective cases, and the stratiform numbers are averaged over the three stratiform cases. Correlation is the Pearson correlation r.

Convective Events			Stratiform Events		
Reference Height z_0 (km)	Equivalent Horizontal Distance Δx (km)	Correlation	Reference Height z_0 (km)	Equivalent Horizontal Distance Δx (km)	Correlation
1.5	2.6* Δz	0.94	1.5	5.5* Δz	0.93

2.0	2.3* Δz	0.95	2.0	6.0* Δz	0.93
2.5	2.7* Δz	0.97	2.5	6.5* Δz	0.91
3.0	2.6* Δz	0.98	3.0	7.2* Δz	0.90
3.5	2.5* Δz	0.96	3.5	7.3* Δz	0.94
4.0	2.9* Δz	0.92	4.0	8.3* Δz	0.96
5.0	2.5* Δz	0.99	5.0	9.6* Δz	0.91
6.0	3.9* Δz	0.95	6.0	7.9* Δz	0.91
7.0	7.3* Δz	0.95	7.0	10.5* Δz	0.99
8.0	5.9* Δz	0.96	8.0	11.7* Δz	0.98

The high correlation values (at least 0.90 for convective and stratiform events) indicate that these simple linear relationships between horizontal and vertical distances are a good approximation. As we would expect based on the example in Fig. 3.10, the multiplicative factors are lower for convective events than for stratiform ones. For example, if we were to replace a pixel at a reference height of 3.0 km with one at 4.0 km (i.e., with Δh = 1.0 km), we would get the same error as we would with a pixel at a horizontal distance of 2.6 km for a convective event, and 7.2 km for a stratiform event—this tendency reflects the stronger horizontal homogeneity of stratiform precipitation. Note also that at the highest reference heights (e.g., greater than 5 km), the multiplicative factors for both convective and stratiform events are higher than those for lower reference heights. This is a reflection of the fact that, at these heights, the vertical gradients near storm top vary so dramatically that any kind of vertical extrapolation is extremely difficult.

The same process has been used in Table A.2 to obtain time equivalences: rather than seeking the lowest variance at a different height level, we determine the lowest variance at a different time. That is, at what horizontal distance is replacement with a pixel at a different time a better option than replacement with a pixel at a different horizontal location? The result is a series of plots similar to Fig. 3.10, and the weighting factors used for calculating equivalent horizontal distances are plotted in Table A.2.

Table A.2: Error equivalence coefficients for time, where Δt (in minutes) is the difference between the time at the reference pixel we wish to replace (t_0) and the time at any potential replacement pixel ($t = t_0 + \Delta t$). These values can be used to calculate the horizontal distances Δx that are approximately equivalent to the time differences Δt in terms of variance. Values are calculated separately for each of the ten possible reference heights. The convective numbers are averaged over all six convective cases, and the stratiform numbers are averaged over the three stratiform cases. Correlation is the Pearson correlation r. Note that the multiplicative factors now have units of (km/min).

Convective Events			Stratiform Events		
Reference Height z_0 (km)	Equivalent Horizontal Distance Δx (km)	Correlation	Reference Height z_0 (km)	Equivalent Horizontal Distance Δx (km)	Correlation
1.5	0.37*Δt	0.99	1.5	0.72*Δt	0.99
2.0	0.35*Δt	0.97	2.0	0.74*Δt	0.99
2.5	0.37*Δt	0.99	2.5	0.76*Δt	0.99
3.0	0.34*Δt	0.99	3.0	0.74*Δt	0.99
3.5	0.35*Δt	0.98	3.5	0.77*Δt	0.98
4.0	0.39*Δt	0.98	4.0	0.77*Δt	0.98
5.0	0.45*Δt	0.98	5.0	0.83*Δt	0.98
6.0	0.57*Δt	0.98	6.0	0.81*Δt	0.98
7.0	0.61*Δt	0.98	7.0	0.93*Δt	0.98
8.0	0.70*Δt	0.97	8.0	1.07*Δt	0.98

As with the vertical pixel replacement in Table A.1, the use of a simple linear relationship between the error for horizontal displacement and the error for temporal displacement is an excellent approximation, with correlation values greater than 0.96 for both convective and stratiform events. Once again, multiplicative factors are lower for convective events than for stratiform ones. As an example, consider a pixel at reference height z_0 = 3.0 km that needs to be replaced. At what point is it better to use a pixel from a different time scan (say, 5 minutes after the reference time) instead of a pixel from the same time scan at a different horizontal distance? Using the numbers in Table A.2, for a convective event, the error for Δt = 5 minutes will be equivalent to the error at a horizontal distance Δx = 0.34*5 = 1.7 km. For a stratiform event, the error Δt = 5 minutes out will be equivalent to that of a horizontal distance Δx = 3.7 km—higher than the convective event, since stratiform events tend to be more horizontally homogeneous.

Appendix B: A brief introduction to ordinary kriging

This appendix will provide a brief introduction to the interpolation method used in the simple and smart ordinary kriging (OK) algorithms of this study. OK has become a staple of the geostatistical toolbox since its origins in the Master's thesis of Danie G. Krige (1951), who applied this method of interpolation and extrapolation as a geostatistical means of ore valuation. Its versatility extends to a wide range of applications, ranging from precipitation mapping (Atkinson and Lloyd 1998), prediction of soil properties (Odeh et al. 1995), mapping contaminated groundwater (Chowdhury et al. 2010), and even non-geophysical applications such as minimizing the torque ripple of a switched reluctance motor (Zhang et al. 2011).

The broad applicability of the OK method is largely due to its simplicity: using a combination of representative pixels weighted using variogram data, missing data can be interpolated or extrapolated. The simple assumptions underlying the method (Wackernagel 1995) are:

- the mean is unknown but constant in the local neighbourhood (in our case, the mean is the actual reflectivity value of a pixel at a specific location and time contaminated by ground clutter); and
- the variogram is known (see Chapter 3 for derivation of reflectivity variograms).

The first step in ordinary kriging, given that the variogram is known, is to determine the weights that will be assigned to each of the pixels involved in the kriging process, so that the value of reflectivity at the pixel in question is given by $Z = \sum_{i=1}^{n} w_i Z_i$, where Z_i are the reflectivity values at a selection i of surrounding pixels. The sum of all weights must be unity for the unbiasedness condition to hold, i.e., $\sum_{i=1}^{n} w_i = 1$. The system of equations to determine the weights in ordinary kriging is simply $\mathbf{W} = S^{-1}R$, with:

- $W = \begin{bmatrix} w_1 \\ \vdots \\ w_n \\ \lambda \end{bmatrix}$, where λ is the Lagrange multiplier used to honor the unbiasedness condition;
- $S = \begin{bmatrix} v(d_{11}) & \cdots & v(d_{1n}) & 1 \\ \vdots & \ddots & \vdots & \vdots \\ v(d_{n1}) & \cdots & v(d_{nn}) & 1 \\ 1 & \cdots & 1 & 0 \end{bmatrix}$, where, for example, $v(d_{1n})$ is the variance at the relative distance between pixel 1 and pixel n (provided by the variograms generated in Chapter 3); and
- $R = \begin{bmatrix} v(d_{1p}) \\ \vdots \\ v(d_{np}) \\ 1 \end{bmatrix}$, where $v(d_{1p})$ is the variance at the relative distance between pixel 1 and the *reference pixel*, i.e., the pixel that is to be replaced.

The preceding equations are for the 2D case where the pixel to be replaced, as well as all its surrounding pixels, are located in the same x-y plane, and where the distance is simply $d_{12} = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$. In Chapter 3, however, we derived variograms in all four dimensions, and we seek pixels that may be offset from the reference pixel in the z- or t-directions. The conversion of the ordinary kriging equation sets from two to four dimensions is trivial: the "distance" between any two pixels is now a distance calculated in four dimensions, using the equivalences between heights, time, and horizontal distance shown in Tables A.1 and A.2 (Appendix A). That is, in order to calculate the equivalent distance between two pixels in four dimensions, we must first convert the height and time differences between them into differences in horizontal distance before proceeding as we normally would to calculate the distance between two pixels in 4D space.

Bibliography

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.

Bordoy, R., J. Bech, T. Rigo, and N. Pineda, 2010: Analysis of a method for radar rainfall estimation considering the freezing level height. *Tethys*, **7**, 25-39.

Chandra, A.S., P. Kollias, S.E. Giangrande, 2010: Long-term observations of the convective boundary layer using insect radar returns at the SGP ARM climate research facility. *J. Climate*, **23**, 5699-5714.

Chilson, P.B., W.F. Frick, J.F. Kelly, K.W. Howard, R.P. Larkin, R.H. Diehl, J.K. Westbrook, T.A. Kelly, T.H. Kunz, 2012: Party cloudy with a chance of migration: Weather, radars, and aeroecology. *Bull. Amer. Meteor. Soc.*, **93**, 669-686.

Chowdhury, M., A. Alouani, and F. Hossain, 2010: Comparison of ordinary kriging and artificial neural network for spatial mapping of arsenic contamination of groundwater. *Stoch. Environ. Res. Risk Assess.*, **24**, 1-7.

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.

Fournier, J.D., 1999: Reflectivity-rainfall rate relationships in operational meteorology. *Local Research Study*, National Weather Service Forecast Office, Tallahassee, FL. (http://www.srh.noaa.gov/tae/?n=research)

Galli, G., 1984: Generation of the Swiss radar composite. Preprints, *22nd Conf. on Radar Meteorology*, Zurich, Amer. Meteor. Soc., 208-211.

Gauthreaux, S.A., Jr., 2006: Bird migration: Methodologies and research trajectories (1945-1955). *Condor*, **98**, 442-453.

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.

Koistinen, J., 1991: Operational correction of radar rainfall errors due to the vertical reflectivity profile. Preprints, *25th Int. Conf. on Radar Meteorology*, Paris, Amer. Meteor. Soc., 91-96.

Krige, D.G., 1951: *A statistical approach to some mine valuations and allied problems at the Witswatersrand* (MS thesis). University of Witswatersrand, 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.

Liu, S., Q. Xu, and P. Zhang, 2005: Identifying Doppler velocity contamination caused by migrating birds. Part II: Bayes identification and probability tests. *J. Atmos. Oceanic Technol.*, **22**, 1114-1121.

Marshall, J.S., W. Hitschfield, and K.L.S. Gunn, 1955: Advances in radar weather. *Adv. Geophys.*, **2**, 1-56.

Martin, W.J., and A. Shapiro, 2007: Discrimination of bird and insect radar echoes in clear air using high-resolution radars. *J. Atmos. Oceanic Technol.*, **24**, 1215-1230.

Odeh, I.O.A., A.B. McBratney, and D.J. Chittleborough, 1995: Further results on prediction of soil properties from terrain attributes: heterotropic cokriging and regression-kriging. *Geoderma*, **67**, 215-226.

Palmer, R.D., D. Bodine, M. Kumjian, B. Cheong, G. Zhang, Q. Cao, H.B. Bluestein, A. Ryzhkov, T.-Y. Yu, and Y. Wang, 2011: Observations of the 10 May 2010 tornado outbreak using OU-PRIME. *Bull. Amer. Meteor. Soc.*, **92**, 871-891.

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.

Scarchilli, G., E. Gorgucci, V. Chandrasekar, T.A. Seliga, 1993: Rainfall estimation using polarimetric techniques at C-band frequencies. *J. Appl. Meteor.*, **32**, 1150-1160.

Uijlenhoet, R., 2001: Raindrop size distributions and radar reflectivity-rain rate relationships for radar hydrology. *Hydrol. Earth Syst. Sci.*, **5**, 615-627.

Wackernagel, H., 1995: Multivariate Statistics—An Introduction with Applications. Springer, Berlin, 389 pp.

Zhang, P., S. Liu, and Q. Xu, 2005: Identifying Doppler velocity contamination caused by migrating birds. Part I: Feature extraction and quantification. *J. Atmos. Oceanic Technol.*, **22**, 1105-1113.

Zhang, Y., B. Xia, D. Xie, and C.S. Koh, 2011: Optimum design of switched reluctance motor to minimize torque ripple using ordinary Kriging model and genetic algorithm. Preprints, *2011 Int. Conf. on Electrical Machines and Systems*, IEEE, 1-4.