A close shave with Occam's razor: managing complexity in a radar rainfall estimation system

Alan Seed and Mark Curtis

[1]*Centre for Australian Weather and Climate Research, Bureau of Meteorology, Melbourne, Australia*

(Dated: 17 July 2014)

Alan Seed

1 Introduction

Advances in computation power have made it possible to explore increasingly complex algorithms to address the vexed question of how to convert radar reflectivity that is observed aloft into an estimate of rain rate on the ground. This year marks my thirtieth anniversary of the first time that I wrote computer software to solve this problem, so it occurs to me that now will be a good occasion to reflect on what I may or may not have learned in the process. The fundamental question that I want to explore is "If some is good, is more better?" In the context of radar rainfall estimation systems, it is instructive to understand both the benefits and the costs of using increasingly complex algorithms.

The Australian Bureau of Meteorology operates a heterogeneous network of 70 weather radars that are located in a wide range of climates (tropics to mid-latitudes, rain forest to desert) and have a wide range of technical capabilities. Understanding the performance of a particular algorithm under increasing levels of noise and in a range of climates is therefore an important issue when considering how to optimize the performance of the radar rainfall estimation system.

The Bureau developed a quantitative radar rainfall estimation (QPE) system Rainfields (Seed et al, 2007). Rainfields 2 became operational in 2008, and testing has recently started for Rainfields3 following a review of the algorithms and system architecture. Rainfields addresses the major components that are required for an operational QPE system and considers partial occultation, clutter identification, corrections for the vertical profile of reflectivity, and real-time adjustment against rain gauge observations (Chumchean et al, 2006). It became apparent during the review of these algorithms that much of the dogma that comes with QPE does not really stand up to close inspection for an operational QPE system.

This paper will use issues that are related to the vertical profile of reflectivity and clutter identification to illustrate some points regarding these fundamental concepts and how they might be handled in a QPE system. A discussion on how the Bureau designed Rainfields3 to account for a heterogeneous network follows, and finally there are some concluding remarks.

2 Managing complexity

2.1 Vertical profile of reflectivity

The impact of melting snow on radar reflectivity observations has been noted since the dawn on radar meteorology e.g. Byers and Coons (1947) and a great deal has been published since then on how to mitigate these effects for quantitative precipitation estimation. The simplest possible option is to ignore the problem entirely and this is not an unreasonable proposition for a radar in the tropics where the altitude of the bright band is around 5000 m ASL. The problem for Australia is that you can almost but not quite entirely get away with this option, and a solution becomes ever more pressing as you go south towards the mid-latitudes in the winter months, for example the bright band is typically below 1500 m ASL in winter for our most southerly radar.

Typically six parameters are needed to describe the mean vertical profile of reflectivity (VPR), a reference reflectivity, the slopes below and above the bright band, the altitude of the bright band and two parameters to describe the shape of the bright band. The Bureau uses the following formulation for the VPR

$$z(h)=\begin{cases} z_0 + a h_b + b(h-h_b) + c\exp\left(-\left|\dfrac{(h-h_b)}{d}\right|^3\right) & \mathrm{h} > \mathrm{h_b} \\ \mathrm{z}_0 + a h_b + c\exp\left(-\left|\dfrac{(h-h_b)}{d}\right|^3\right) & \mathrm{h} < \mathrm{h_b} \end{cases}$$

where z_0 is the reflectivity at the ground, h_b is the height of the bright band, a and b are the slopes of the VPR below and above h_b respectively, c is the intensity and d is the width (Snow et al 2012).

The radar reflectivity that is observed aloft is extrapolated onto the ground using the anisotropic Ordinary Kriging algorithm of Seed and Pegram (2001). This algorithm uses Ordinary Kriging to extrapolate a set of the the nearest non-clutter radar reflectivity bins onto the ground. The semi variogram accounts for the fact that the vertical structure of radar reflectivity is not the same as the horizontal structure and removes the bias that is caused by the VPR. The mean square error for the extrapolation is calculated and used to construct a quality index map that is used when combining the QPE products from several radars into a mosaic.

Figures 1 and 2 show a daily rainfall accumulation of radar rainfall estimates with and without a VPR correction. It is clear that in this case the positive bias close to the radar has been reduced, but the negative bias at far range over the land has not been changed.

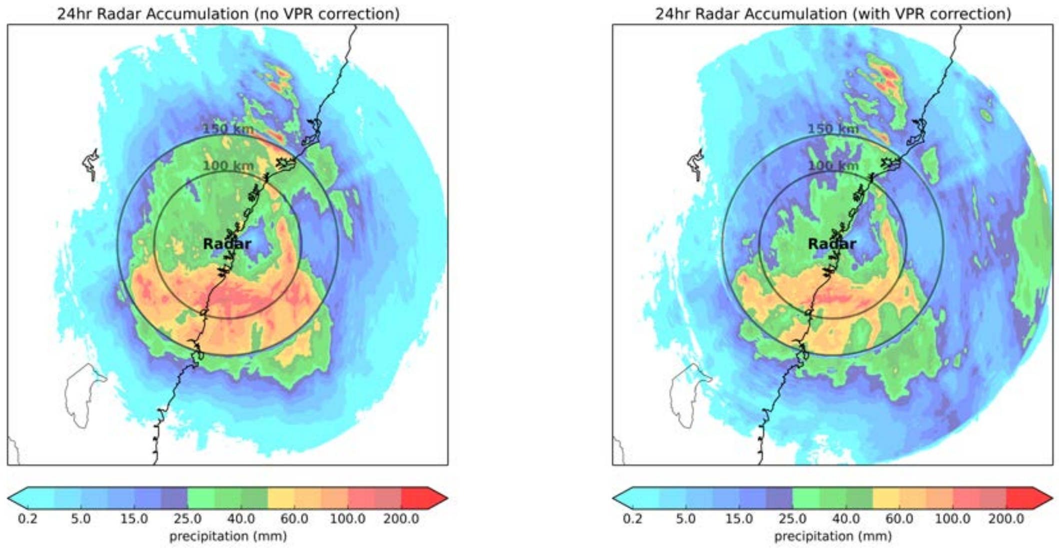

Figure 1 Example of a daily accumulation of radar rainfall estimates with and without VPR correction

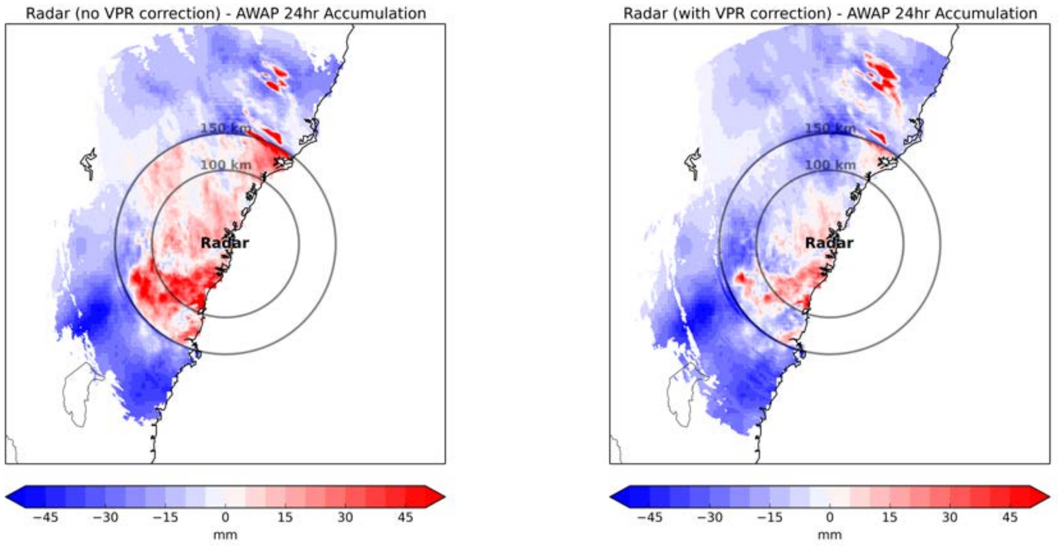

Figure 2 Difference between daily accumulations of radar rainfall estimates with and without VPR correction and a gridded gauge ground "truth"

If we assume for the moment that the mean vertical profile of reflectivity is a valid construct, then two practical problems become apparent. Firstly there is the problem of using radar observations to estimate the parameters in real-time. Typically this is done using observations that are "close" to the radar, which usually means the range where the diameter of the radar beam is significantly narrower than the height of the bright band and where the base scan is below the bright band, since only the slope of the VPR above the bright band can be reliably estimated at far range. There is also a problem of accurately determining the slope of the VPR under the bright band since it is not always possible to get clean radar reflectivity observations at low elevations due to ground clutter. An optimal solution is to simplify the model to using a forecast of the wet bulb freezing level from a Numerical Weather Prediction model and having fixed or slowly varying values for the other parameters when it is determined that the errors from the fully parameterized model will exceed those from a constrained model.

Secondly, there is the problem of only applying the correction to the cases where there is in fact a bright band since errors of the order of 3-5 dBZ can arise when "correcting" a profile of a convective cell that does not have a bright band. Figure 3 shows an example of a situation where small winter convective showers were moving rapidly away from the radar. The apparent non-vertical structures are caused by the motion of the cells between the elevation scans in the volume. This illustrates that the vertical structure of radar reflectivity above a location on the ground is very noisy and the mean VPR only explains a small fraction of the observed variability. Therefore it is very difficult to detect the presence of a bright band reliably and the correction for the bright band is necessarily very noisy as a result.

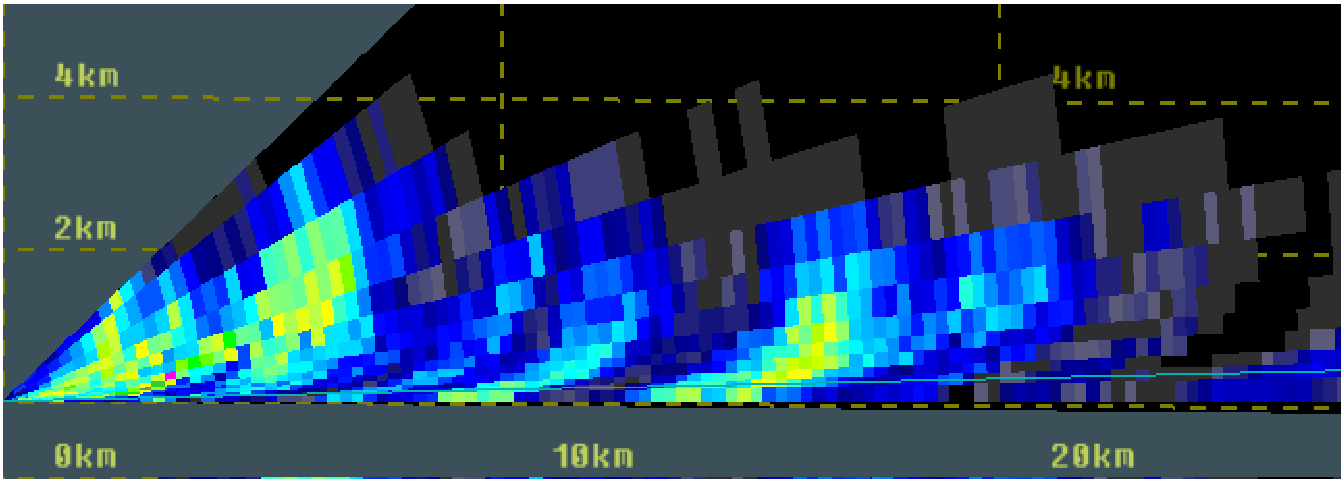

Figure 3 Example of a vertical cross section through a volume scan for the Brisbane Radar

So in conclusion, it is very difficult to estimate the parameters of a VPR model reliably in real-time, the VPR only explains a small fraction of the observed variance in the vertical, and the VPR for a convective cell is significantly different from that of an area of stratiform rain. Despite all this, a VPR correction clearly improves the quality of long-duration accumulations of radar rainfall estimates and is worth doing, but what is an appropriate level of complexity in the model that is used to correct for the VPR?

2.2 Clutter identification

Peter et al. (2013) developed a Naïve Bayes Classifier (NBC) that was developed by to discriminate between anomalous propagation, sea clutter, and precipitation. This was further modified by Rennie et al. 2014 to include the classification of clear-air targets that provide radial velocity observations that are suitable for assimilating into a NWP model. The feature fields used in the classification included measures of texture of reflectivity and radial velocity, echo top height, reflectivity, and spectrum width. The Australian network does not have an operational dual-polarametric radar, and the older radars are reflectivity only.

Table 1. List of the classes used for manual classification: Class number, class name and class abbreviations as used for figure labels, etc.

Class number	Class	Class abbrev.
1	Convective precipitation	con
2	Shallow convective precipitation	shc
3	Stratiform precipitation	str
4	Insects	ins
5	Smoke (bushfires)	smk
6	Chaff	chf
7	Macro-aerofauna (birds/bats)	brd
8	Permanent ground clutter	pe
9	AP ground clutter	gc
10	AP sea clutter	ap
11	Side-lobe sea clutter	sl
12	Second trip echo	2tp

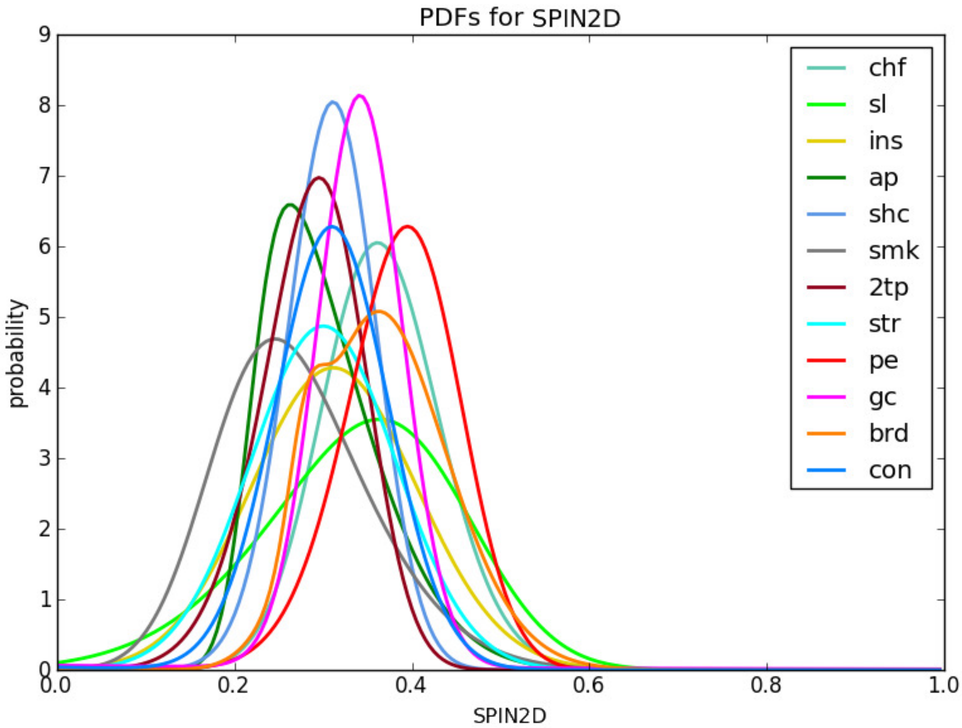

Figure 4 Probability distribution of the spin feature field for the various classes. Rennie et al, 2014.

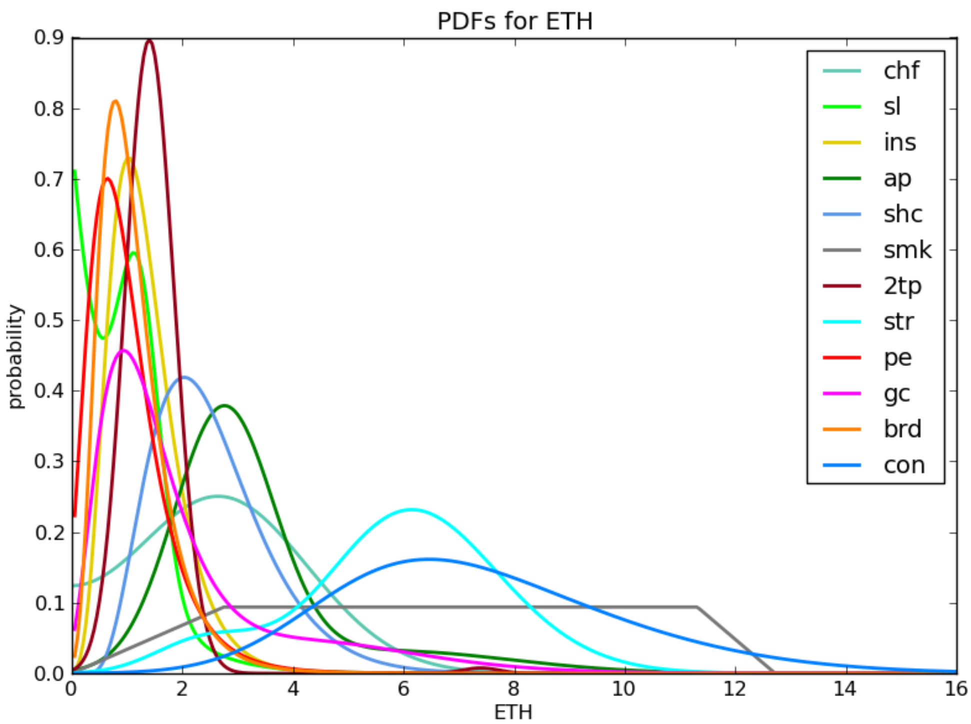

Figure 5 Probability distribution of echo top height for the various classes. Rennie et al, 2014

Table 6. Comparison of the Bayesian classification method with an existing threshold-based method, discriminating between precipitation and non-precipitation.

		Bayesian				Thresholds	
		Automatic				Automatic	
		precip.	non			precip.	non
Manual	precip.	94.4%	5.6%	Manual	precip.	90.3%	9.7%
	non	23.5%	76.5%		non	46.7%	53.3%

Verification of the Bayes classification scheme shows that it is an improvement on a very simple scheme for QPE that uses vertical gradient and echo top height thresholds to identify rain and reject all the other possibilities. The classification scheme was optimized to provide clean radial velocity observations to a data assimilation scheme, not to provide reflectivity to a QPE system. The impact of this rather subtle difference can be seen in Figure 6 where stratiform rain has been identified as "smoke".

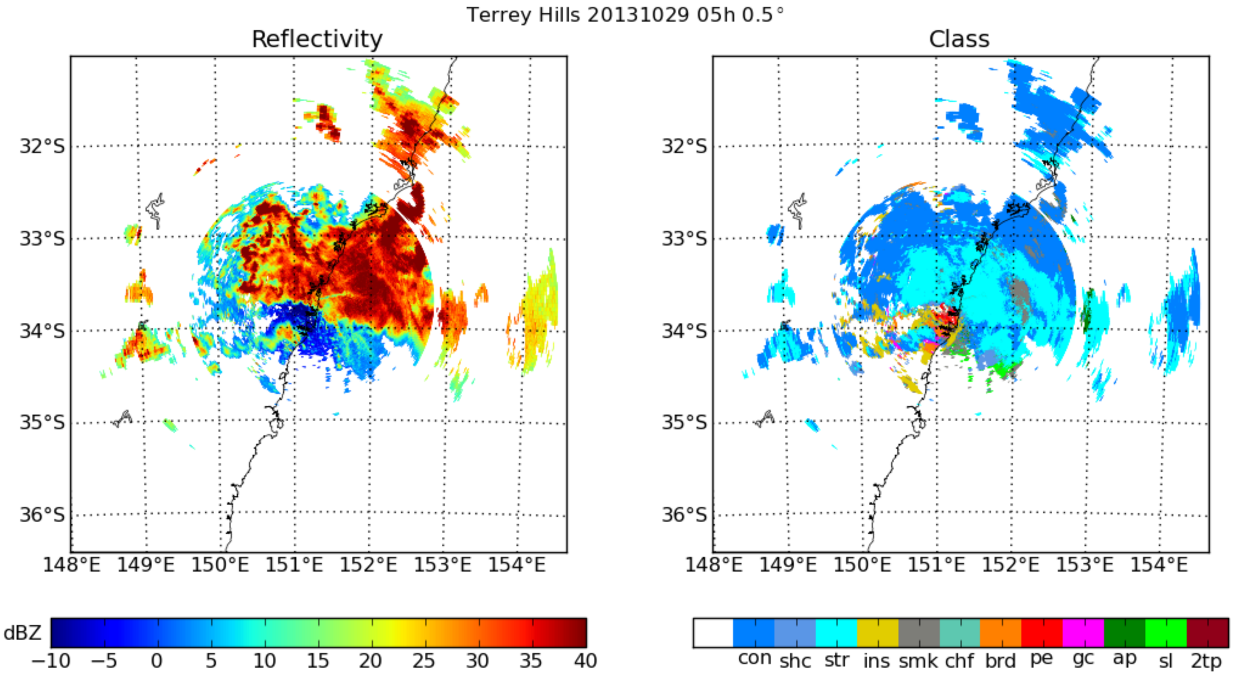

Figure 6 Example of a radar reflectivity field and associated target identification, Rennie et al 2014.

This illustrates two points. Firstly, that it is very difficult to optimize a quality control system for both radar QPE and radial velocity applications and either a compromise or two parallel systems is required. The cost of managing two operational quality control systems is very high in terms of system complexity and human resources that are needed to manage and calibrate the two systems, so a compromise solution is required. Secondly, from a QPE perspective there is a modest increase in the average performance of the quality control system for rainfall at the expense of a very significant effort. The miss-classification of rain as say smoke does more harm to the QPE system than simply missing out on some shallow rainfall from time to time. Therefore the cost functions that are used when calibrating a radar quality control system need to reflect occasional errors that have a large impact on a specific application.

3 Building a configurable QPE system

When developing a QPE system it can be tempting to give almost exclusive attention to the implementation of the various quality control and product generation algorithms. This work is of course critical, however it is by no means the only functional area to which focus must be applied. There are three distinct functional blocks, or work-flows, for which a modern QPE system must take full account. These are the areas of product life-cycle, production management, and product generation.

The product life-cycle function is responsible for storage, retrieval, searching and eventual archive or removal of the various QPE products. This function acts as the interface to both internal and external clients. The production management function handles the computational resources needed for the timely generation of products including triggering, monitoring and dynamic scaling of the system. Finally, the product generation function is responsible for implementation of all the scientific and engineering processes needed to take raw radar and rain gauge observations and produce high quality precipitation estimates.

Each of these work-flows must scale independently as the system grows, and for this reason they should be implemented with a high degree of independence. Product life-cycle functions must continuously scale as new products need to be stored and managed alongside those which are retained. Production management must scale up as the network of radars and rain gauges expands or contracts. Product generation must scale up as new technologies find adoption within the network, requiring continuous evaluation and innovation of quality control measures. In Rainfields 3 each of these areas is implemented separately allowing the system to cleanly adapt as the situation in which it is deployed evolves.

In addition to changes over time, the system must also cope with a diverse range of radar technologies situated within a wide range of climate conditions. In the Australian weather radar network there exists a mixture of C and S band radars, using 1, 1.5 or 2 degree antennas with non-Doppler, Doppler or dual-polarization receivers, scattered over approximately 30 degrees of latitude. The network is far too diverse to allow for a ‘one size fits all’ approach to the QPE algorithms, yet the sheer number of combinations represented makes building dedicated software for each situation infeasible.

The approach used within Rainfields 3 is to split the QC and QPE processes into a collection of algorithmic building blocks which may be assembled in any combination and order to suit an individual radar. Each building block, known as a ‘stage’, is responsible for a single clearly specified part of the quality control process. A single stage may implement a reasonably complex function, such as correcting for partial beam blocking, or a very simple function such as multiplying a field by a scalar. Inputs for one stage are normally driven by the outputs from a previous stage, and in this way multiple stages are chained together to form a cohesive QC/QPE process customized to the needs of a particular radar. In the current Rainfields 3 deployment a typical radar processing chain is comprised of over 50 stages.

Although each stage is implemented within software, the configuration of the processing chain (including the linking of inputs and outputs, and setting of various adjustable parameters exposed by each stage) is read from an XML file. This allows tracking and change management of the quality control and precipitation estimation processes independently from that of the software itself where updates are often more to do with quality of implementation than changes to the processing approach.

To prevent the number of unique processing chain configurations required from suffering the same combinational explosion as would dedicated per-radar software, each stage provides default values for its configuration parameters. These defaults are maintained and updated as a matter of course when quality control techniques improve. In this way all radar configurations may benefit from improved parameter values without manual intervention.

For radars where a default configuration parameter is unsuitable it may be overwritten on a per-radar basis. Such updates are time stamped and tracked by the system, allowing the later regeneration of very long periods of products by seamlessly tracking the original changes in site configurations.

4 Conclusions

It is evident that there is a steady increase in the complexity of algorithms for radar rainfall estimation in the published literature. This is a result of the increasing power of computers to support ever more complex algorithms, and the natural bias towards complexity in our approach to research. "If some is good, then more is not better" when it comes to building a robust operational radar rainfall system, and Occam's razor, "Everything should be as simple as possible, but not simpler", is an important design concept that needs to be applied at each stage in the development of the QPE system. Gratuitous complexity comes at the cost of increased difficulty in calibrating the system and difficulty in understanding just how the system will perform in sub-optimal situations.

References

Byers H.R., and Coons R.D. (1947), The "Bright Line" in radar cloud echoes and its probable explanation, J. Meteorol., 4, 75 – 81.

Chumchean S., Sharma A., and Seed A.W. (2006), An integrated approach to error correction for real-time radar rainfall estimation. J. Atmos. and Oceanic Tech., 23, 67 – 79.

Peter J. R., Seed A.W., and Steinle P. (2013), Application of a Bayesian Classifier of Anomalous Propagation to Single-Polarization Radar Reflectivity Data. J. Atmos. Oceanic. Tech., 30, 1985-2005.

Rennie S.J., Curtis M., Peter J.R., Seed A.W., and Steinle P.J. (2014), Training the Ancilla Naïve Bayes classification system to provide quality control for weather radar data. CAWCR Technical Report No. 073, June 2014. http://cawcr.gov.au/publications/technicalreports/CTR_073.pdf

Seed A.W., Duthie E., and Chumchean S. (2007), Rainfields: The Australian Bureau of Meteorology system for quantitative precipitation estimation. Proceedings of the 33rd AMS Conference on Radar Meteorology.

Seed A.W., and Pegram G.G.S. (2001), Using Kriging to infill gaps in radar data due to ground clutter in real-time. Proceedings of the 5th International Symposium on Hydrological Applications of Weather Radar, Kyoto, Japan, 19-22 November 2001, 73-78.

Snow K., Seed A.W., and Takacs G. (2012), Determining the vertical profile of reflectivity using radar observations at long range. Proceedings of the Weather Radar and Hydrology Conference, Exeter, U.K. 2012, IAHS Publ. 251. 69 - 74.