

GRAS SAF Report 01
Ref: SAF/GRAS/METO/REP/GSR/001
Web: www.grassaf.org
Date: 1 October 2008

The EUMETSAT
Network of
Satellite Application
Facilities



GRAS SAF Report 01

Mono-dimensional Data Thinning for GPS Radio Occultations

Carlo Buontempo

Met Office, UK

Document Author Table

	Name	Function	Date	Comments
Prepared by:	C. Buontempo	GRAS SAF Project Team	1 October 2008	
Reviewed by:	D. Offiler	GRAS SAF Project Team	1 March 2007	
Approved by:	K.B. Lauritsen	GRAS SAF Project Manager	17 October 2008	

Document Change Record

Issue/Revision	Date	By	Description
Version 0.1	01 Mar 2006	CB	Preliminary release for I-RR
Version 0.2	25 Sep 2006	CB	Added detail to SG algorithm
Version 0.3	27 Nov 2006	DO	Add implementation chapter. General tidy up
Version 1.0	01 Mar 2007	DO	Update Implementation chapter. First release version
Version 1.1	01 Oct 2008	HL	Included report as part of GSR series. Removed implementation chapter - to be part of GSR 08

GRAS SAF Project

The GRAS SAF is a EUMETSAT-funded project responsible for operational processing of GRAS radio occultation data from the Metop satellites. The GRAS SAF delivers bending angle, refractivity, temperature, pressure, and humidity profiles in near-real time and offline for NWP and climate users. The offline profiles are further processed into climate products consisting of gridded monthly zonal means of bending angle, refractivity, temperature, humidity, and geopotential heights together with error descriptions.

The GRAS SAF also maintains the Radio Occultation Processing Package (ROPP) which contains software modules that will aid users wishing to process, quality-control and assimilate radio occultation data from any radio occultation mission into NWP and other models.

The GRAS SAF Leading Entity is the Danish Meteorological Institute (DMI), with Cooperating Entities: i) European Centre for Medium-Range Weather Forecasts (ECMWF) in Reading, United Kingdom, ii) Institut D'Estudis Espacials de Catalunya (IEEC) in Barcelona, Spain, and iii) Met Office in Exeter, United Kingdom. To get access to our products or to read more about the project please go to <http://www.grassaf.org>.

Contents

1 Introduction	2
1.1 Requirements	2
1.2 Aim of the document	2
Bibliography	4
2 Methods for mono-dimensional data thinning	5
2.1 Vertical resolution: constrain due geometry and noise	5
2.1.1 Geometric optics	5
2.1.2 A limit to vertical resolution due to the measurement error	5
2.2 Thinning through sampling	6
2.3 Thinning strategies: advantages and disadvantages	8
2.3.1 Constant sampling	8
2.3.2 Incremental sampling	8
2.3.3 Adaptive sampling	10
2.4 Thinning through interpolation	13
2.4.1 Linear interpolation	13
2.4.2 Logarithmic interpolation	14
2.4.3 Polynomial interpolation	14
2.4.4 Spline interpolation	14
2.5 Smoothing	16
2.5.1 SG filter: a way to thin and smooth the data coherently	16
2.6 Conclusion	17
Bibliography	19
A Glossary	20
Bibliography	22
B Thinning Levels	23

1 Introduction

Thinning is generally part of a long sequence of operations performed on data before it becomes useful. Thinning aims to reduce the amount of data without reducing the information content. In the case of radio occultation the preprocessing involve also: a smoothing and a statistical optimisation. The latter has been thoughtfully analysed by Healy (2001) and will be not further discussed in this report which is devoted to the analysis of the thinning strategies, of their differences and of their specific advantages. As we will see later, smoothing is neither necessary nor sufficient to thin the data. The two processes are theoretically separate and are usually joined together only for practical purposes.

1.1 Requirements

The implemented thinner must:

- reduce data volume with minimal loss of information content
- not introduce any artifact such as vertical correlations
- allow BUFR encoding such that message lengths are compatible with dissemination over the GTS, In practice this means messages less than 15,000 bytes
- run quickly and add no significant overhead to the encoding itself — i.e. proces in a profile in seconds (or less) rather than minutes
- output on fixed levels, preferably those defined in Chapter B
- be flexible enough that different sets of arbitrary levels can be user-defined
- implement alternative, user-selectable methods

1.2 Aim of the document

In the case of GPS radio occultation (RO), data tends to be oversampled in respect of the actual resolution and information content of the dataset. Although this approach offers several technical advantages, it also unnecessarily increases the bandwidths for transmission,

storage volumes and the computational needs of NWP centres when assimilating the RO data.

This GRAS SAF Report aims to present and discuss some possible ways of reducing the amount of data without losing information. This operation is normally referred to as thinning. Because GPSRO are substantially monodimensional (profiles) only 1D data thinning algorithms will be presented here. The multidimensional data thinning, crucial for other satellite application (such as infrared imagers and microwave profilers) will not be analysed.

Bibliography

Healy, S., Smoothing radio occultation bending angles above 40 km, *Ann. Geophys.*, 19, 459–468, 2001.

2 Methods for mono-dimensional data thinning

2.1 Vertical resolution: constrain due geometry and noise

2.1.1 Geometric optics

A typical GPS bending angle profile contains about 5000 points. This could give an erroneous impression on the vertical resolution: $\frac{\sim 60km}{\sim 5000points} \approx 30metres$, whereas the effective resolution of a GPSRO profile is much larger. The first limitation in that sense comes from the diffraction of the antenna. There is theoretical limit to the vertical resolution of a geometric optic retrieval. This is defined by the diameter of the first Fresnel zone (Hinson and Magalhaes, 1991) which is a way of identifying the typical size of the electromagnetic beam. The concept of Fresnel zone is crucial for telecommunication since the presence of an obstacle in the 1st zone can significantly reduce the quality of the transmission. For that reason the first Fresnel's zone is usually defined as the first of an infinite number of a concentric ellipsoids, surrounding the main axis of propagation of the radio signal from where a secondary wave, generated by an obstacle, generates constructive interference at the receiver.

The radius of the Nth Fresnel zone is given by:

$$R_N = \sqrt{\frac{N \times \lambda \times D_1 \times D_2}{D_1 + D_2}} \quad (2.1)$$

Where N is the order of diffraction (the number of Fresnel zone), λ is the wavelength of the signal (~ 20 cm for GPS) and D_1 and D_2 are the relative distances to the two satellites.

Using the typical values of a GPS occultation we have that the first Fresnel's zone reduces the effective resolution to few hundreds meters (Kursinski et al., 1997).

2.1.2 A limit to vertical resolution due to the measurement error

Sampling each Fresnel diameter with few a data points would be more than adequate to represent all the physically available information in an RO profile, but that is not the only limitation to the effective resolution of the profile. Marquardt and Healy (Marquardt and Healy, 2005) have recently shown that it is the noise in RO data that limits the vertical resolution.

Furthermore the recent development of more elaborated inversion algorithms has made a sub-fresnel vertical resolution theoretically achievable (Jensen et al., 2003; Sokolovskiy, 2001; Gorbunov et al., 2000)

The problem of data thinning is strongly related to the concept of independent pieces of information as defined in Rodgers (Rodgers, 2000). The basic idea is that only the quantities whose internal variability is larger than the respective measurement error covariance can be physically measured. This defines the effective rank of the problem. Usually the natural variability of the measurement vector is expressed in terms of prior covariance. In more precise mathematical terms it is possible to say that: the number of independent measurements made to better than measurement error, which defines the effective rank of the problem, is equal to the number of singular values greater than one of

$$S_{\epsilon}^{-\frac{1}{2}} K S_a^{-\frac{1}{2}} \quad (2.2)$$

Where S_{ϵ} is the measurement covariance and S_a is the background covariance and K is the forward model. Marquardt (Marquardt and Healy, 2005) has established that a typical bending angle profile does not contain more than ~ 250 independent pieces of information. This represents the number we should keep in mind for the rest of the discussion. The required set of impact heights used with the ROPP thinning algorithms is given in Chapter B

2.2 Thinning through sampling

The most immediate reason to thin the data is to reduce the cost of processing, storage and transmission: the use of a lower sampling rate usually results in a cheaper implementation. Furthermore from the previous section it is clear that it is of no benefit to process more points than the number of independent pieces of information present in the profile.

Probably the easiest way to reduce the number of data points is to create a sub-sample of the original dataset. This could be implemented in different ways. The simplest is to reduce the size of dataset by selecting (or rejecting) one point every N . We refer to this procedure as constant sampling since the sampling rate remains constant independent of the intrinsic variability of the profile (see the Glossary). An example of the behaviour of the **Constant Sampling (CS)** for different sampling rates is reported in Figure 2.1

In general it would be desirable to increase the sampling rate in the region where most of the information is concentrated. If for instance a large portion of our unthinned dataset lays on a straight line segment [A,B] two points are enough to describe the profile completely. In the same way a parabola fits the data perfectly, only three data point are necessary and so on. Unfortunately we do not know in advance whether the profile will be describable by a

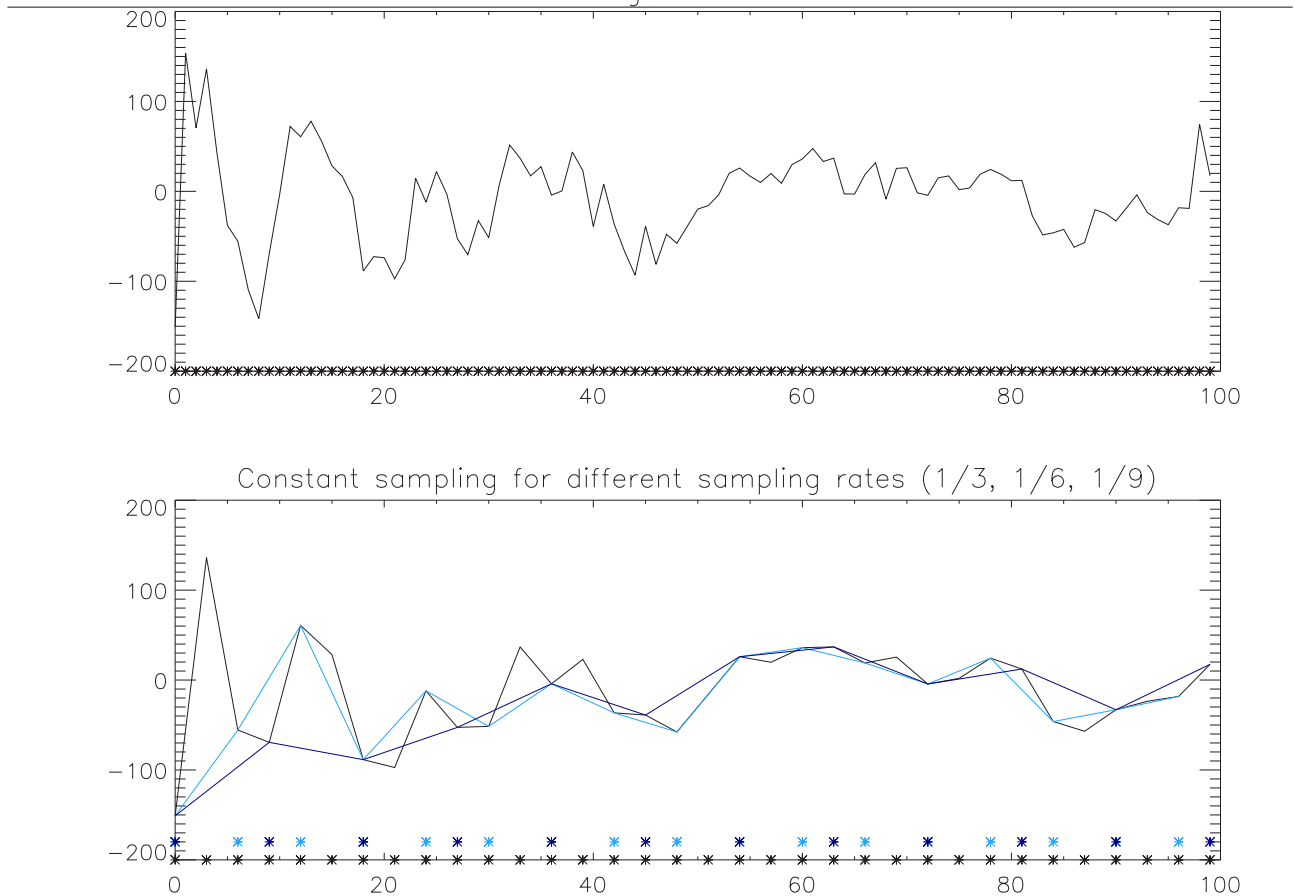


Figure 2.1: The reconstructed profile for different sampling rate: the stars indicate the position of the sampling

first, second or higher order polynomial for a part of it and then we have to rely on different technique to enhance the sampling in the region of interest. The next section describes in more detail some potential thinning strategies, their advantages, their disadvantages and their possible implementation.

2.3 Thinning strategies: advantages and disadvantages

2.3.1 Constant sampling

The zero order approximation, in order to improve the performances of the constant sampling algorithm, is to have a sampling rate which varies from one dataset to another accordingly to its intrinsic variability. This means to increase the sampling for profiles exhibiting large variability and reducing it for more uniform profiles. This has been achieved analysing the lagged autocorrelation of the first derivative of the dataset (e.g. Figure 2.2). In particular the optimal sampling rate has been defined as the lag value for which the autocorrelation falls below a fixed threshold (0.2 in the actual implementation). This approach has been called **Optimal Constant Sampling (OCS)**.

Advantages

- OCS is extremely inexpensive to run
- The desired amount of thinning is easy to set

Disadvantages

- The procedure samples with the same sampling rate independently of the presence of structure in the data
- Areas characterised by large peaks are sampled in much the same way used for more uniform regions
- The thinned profile is produced on verticals level which vary from one case to another

2.3.2 Incremental sampling

OCS is indeed a better approach than the fixed sampling there is still space for improvements. The autocorrelation is a global evaluation of the statistical properties of the dataset. Any attempt to derive the sampling rate from a global measure of variability will end up in oversampling the more uniform areas and undersampling those characterised by an higher variability. This tends to be a problem in the lowest part of the profile which is generally characterised by large gradients in both temperature and humidity and which consequently tends to exhibit larger fluctuation.

It is clear both CS and OCS perform particularly poorly in the area where the original data shows the highest variability, (Figure 2.1). The second order approximation is one-dimensional data thinning is to relate the sampling rate to a local measurement of variability.

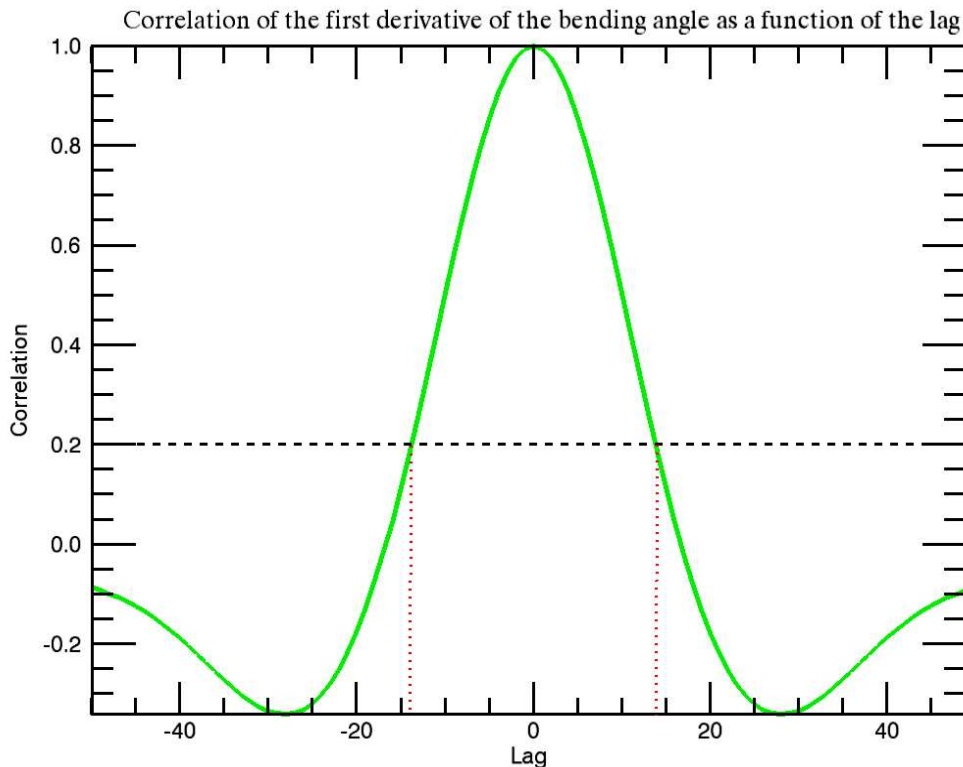


Figure 2.2: The correlation of the first derivative of a bending angle profile. The red dotted line indicate a possible threshold for the sampling step.

An obvious quantity to look at is the second derivative of the data since its absolute value reaches an extreme (either a maximum or a minimum) on the spikes. This approach, that we have called **Incremental Sampling (IS)**, can be conceptually divided into two steps. First a function of the second derivative of the data in respect to time is calculated, then, using a threshold for the value of this function, the algorithm evaluates whether to increase or decrease the sampling rate. There are many possible methods that can be used to implement IS. For the test reported in here we have chosen the average, over 10-point windows, of the absolute value of the second derivative. As shown in Figure 2.3, IS improves the resolution in areas where most of the structure is concentrated; this improves the reconstruction of the profile seen in Figure 2.4.

Advantages

- Preserves resolution in stuctually-variable areas
- Although computationally more expensive than CS, IS is still relatively cheap.

Disadvantages

- As for the OCS the output is not defined on a constant set of altitude
- It is harder to define a rejection rate

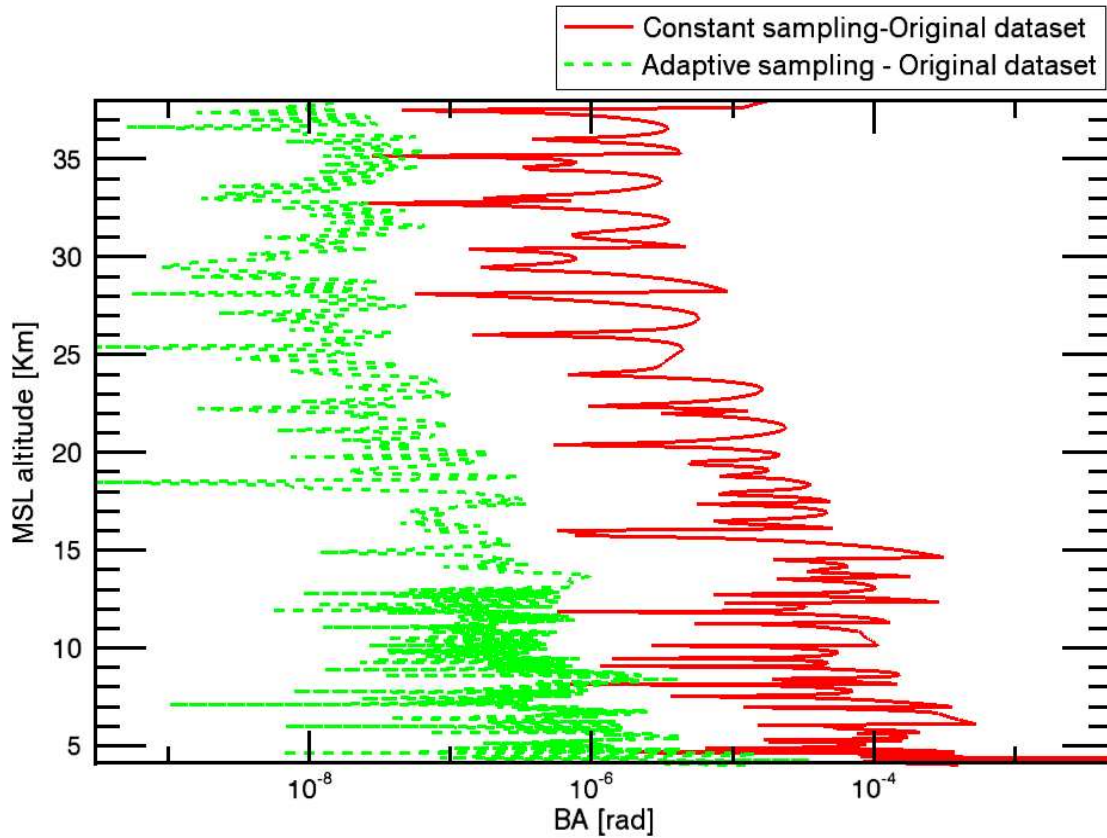


Figure 2.3: A bending angle profile sampled through a constant and an incremental sampling procedure

2.3.3 Adaptive sampling

The aim of any sampling procedure is to increase the sampling in regions where most of the information is supposed to be. If we do not want to use our previous knowledge on the possible structure of the profile the only possibility we have is to oversample in regions where the profile shows high variability. In the previous example (IS) this was achieved by increasing the sampling rate where a monotonic function of the second derivative was exceeded a threshold but the same result can be achieved using other approaches.

The approach to mono-variate data thinning developed by Dyn (Dyn et al., 2000) has been followed to create an **Adaptive Sampling (AS)** procedure. In this case the original profile is divided into sub-samples from which the points are removed one by one. The first points to be removed are those less significant in some way. In the implementation that we have tested for this report, the significance of each point has been related to their unique contribution to the value of the integral. For each sub-sample the complete integral is compared to the integral of a different subsample from which a data point is selectively eliminated. Those data points whose removal least alters the overall integral have been removed. In more mathematical terms, following the notation adopted by Dyn, if we have

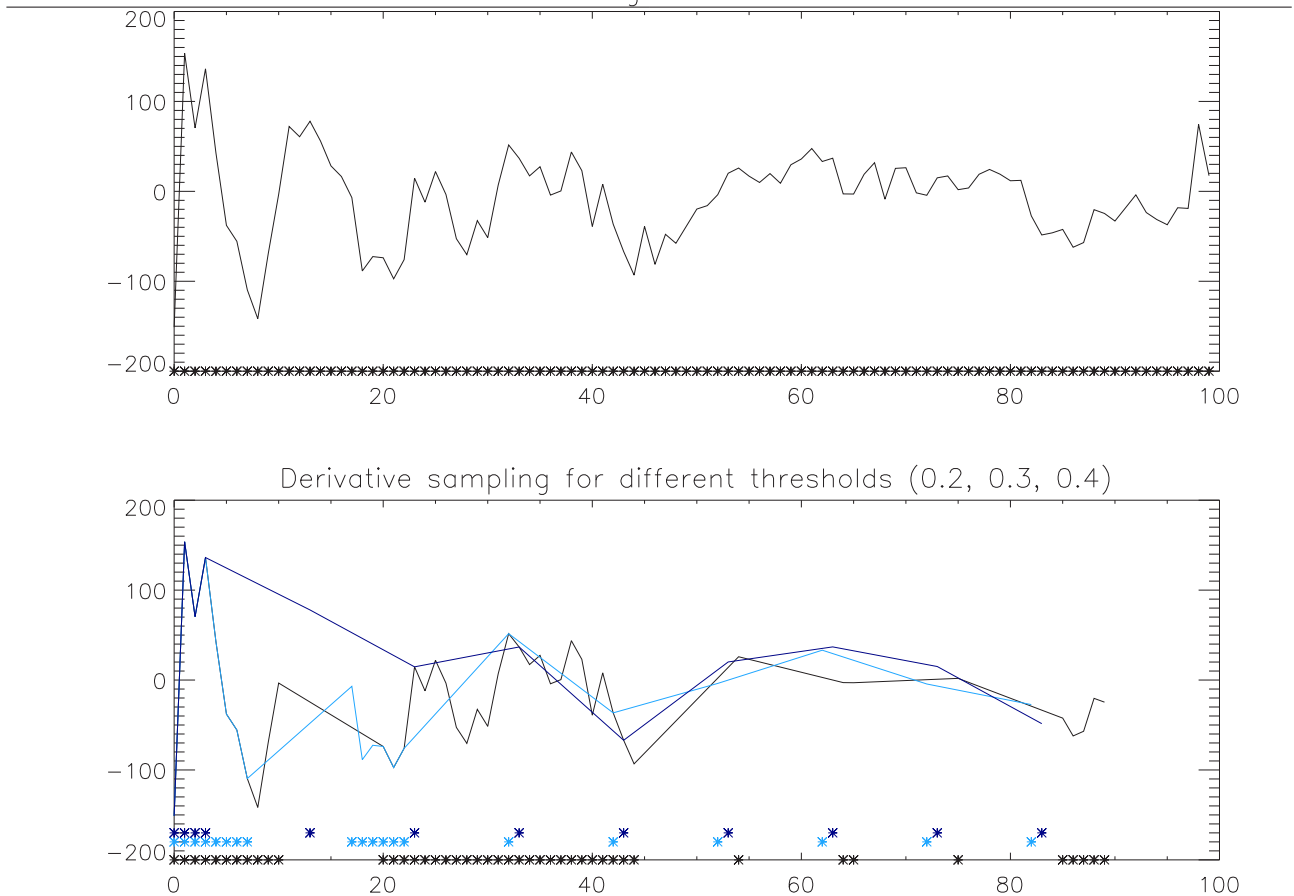


Figure 2.4: The dataset reconstructed using the incremental sampling procedure for different values of the threshold

an interval $[a, b] \in \mathfrak{X}$ and $X = (x_1, \dots, x_N)$ is a given sequence of points in $[a, b]$ in ascending order: $a = x_1 < x_2 < \dots < x_N < b$ and suppose that the function $f \in \mathfrak{X}$ is sampled at those points, the problem of thinning could be expressed as the problem of finding a subset of X $Y = (y_1, \dots, y_n)$ with $n < N$ such that the linear interpolant $L(f, Y)$ to the data is close to the given data in the sense that the error $E(Y; X; f) = \max \|L(f, Y)(x) - F(x)\|$ is small compared to the error of other sub-samples of X of the same size. In our implementation we have adopted the integral as a general interpolant for the sample. This technique has two advantages: it is easy to set the thinning fraction (number of point to throw away) and it is computationally inexpensive. An example of the result obtained with this approach is shown in Figure 2.5.

Advantages

- Although computationally far more expensive than CS, AS is still quite efficient
- The sampling rate changes from on dataset to another
- The sampling rate changes within the same dataset from region to region according to the profile variability

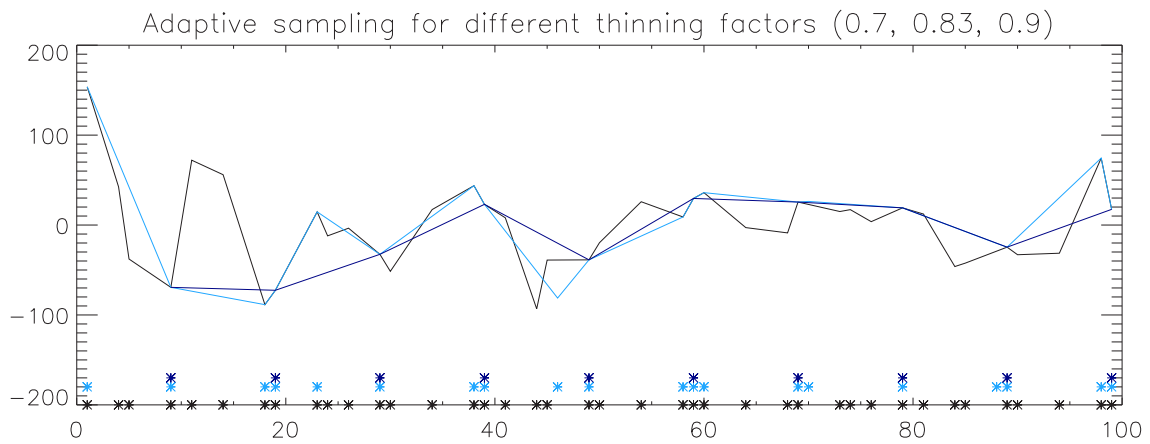



Figure 2.5: The dataset reconstructed using the adaptive sampling procedure for different values of the rejection fraction

Disadvantages

- As in the previous cases the outputs are defined on a set of altitudes which varies from one profile to another.

2.4 Thinning through interpolation

So far we have described only thinning algorithms based on data sampling. This is theoretically the cleanest way to obtain a reduction in the data since the values of the "observed points" are not altered. Thinning can also be achieved by performing an interpolation onto a specific sub-set of altitudes. The Collins English Dictionary has three different definition of interpolation:

1. *to insert between or among others*
2. *to change by putting in new material*
3. *to estimate a missing value by taking an average of known values at neighbouring points.*

Although most of the readers would agree in finding the third answer to be the most appropriate, in a sense all of them are true. Interpolation is a technique which constructs new data points from a discrete set of known data points. Although in most cases this is a safe operation, new data are add where there was none before. Interpolation has two main advantage in respect to sampling: the output can be obtained on a defined set of levels and it easier to implement a pre-smoothing procedure which can reduce random noise in the original data.

2.4.1 Linear interpolation

Probably the simplest method is linear interpolation. In this case the value of the function f in a point $x \in [a, b]$ is, assuming that $(a; b)$ are sampled points,

$$f(x) = \frac{x - x_b}{x_a - x_b} \times y_a + \frac{x - x_a}{x_b - x_a} \times y_b \quad (2.3)$$

Linear interpolation is, in general, not very accurate. Furthermore it is not differentiable since its first derivative is discontinuous. It is possible to demonstrate that the error of a linear interpolation is proportional to the square of the sampling step. Most of the other interpolation schemes have errors proportional to higher powers of the sampling interval.

Advantages

- Although more expensive the CS, it is less expensive then the other kinds of interpolation
- The output is defined on a fixed set of altitudes

Disadvantages

- It is not very accurate
- it is not differentiable

2.4.2 Logarithmic interpolation

Logarithmic interpolation is basically the same as linear interpolation except that the values are first converted to their logarithm, linearly interpolated, and the interpolated points re-converted to their natural values. This method is more accurate for quantities which vary (on the average) logarithmically with height — e.g. bending angle. The advantages and disadvantages for linear interpolation also apply here.

2.4.3 Polynomial interpolation

Polynomial interpolation is a generalisation of the linear interpolation where the linear interpolant is replaced by a polynomial of higher degree. It is possible to demonstrate (The Fundamental Theorem of Algebra) that having n different data points it is possible to find a $(n-1)$ degree polynomial passing through all of them. One of the biggest disadvantage of polynomial interpolation is that is computationally expensive, especially for large datasets.

Advantages

- Accurate interpolation
- The polynomial interpolation is differentiable everywhere (smooth)

Disadvantages

- It can be very expensive for large datasets
- can produce artifacts around the maxima (Runge's phenomenon)

2.4.4 Spline interpolation

Spline interpolation is piecewise polynomial interpolation. Instead of looking for a single polynomial which fits the data exactly, the original dataset is divided into parts for each sub-sample and a polynomial interpolation is performed using a mean square fit.

Advantages

- As for the full polynomial, interpolation is smooth and accurate
- it is computationally much cheaper than full polynomial interpolation
- Doesn't suffer the Runge's phenomenon

Disadvantages

-
- The accuracy of the interpolation depends on the degree of the polynomial adopted.

2.5 Smoothing

The basic assumption behind any smoothing procedure is that the signal we want to analyse is slowly varying in respect to a random noise which is supposed to affect it. In that sense, assuming that the data have not been optimally thinned already (in an information content sense) there is some redundancy between adjacent points. Any procedure which substitutes at each point a sort of local average performs a smoothing of the original dataset. This can be expressed as a digital filter whose general form can be written as

$$g_i = \sum_{n=lmost}^{rmost} c_n f_{i+n} \quad (2.4)$$

where $lmost$ and $rmost$ represent the left most and right most point of the averaging window and c_n are relative weights for each point.

One of the most popular and simple approaches to data smoothing is the running or 'boxcar' average. In this case the function adopted is the average of the n point inside the window. This means that all c_n defined in equation 2.4 have all the same value (unity). This is a clear disadvantage in the region where the signal exhibits large peaks since although the first moment (the integral) of the signal remains unchanged, the second moment (the width of the peak) is affected. Running averages always create peaks which are broader and less intense than the original dataset.

2.5.1 SG filter: a way to thin and smooth the data coherently

In general, smoothing through a function which preserves the higher moment of the distribution (in respect to the typical size of the feature we are interested in) is highly desirable. In their analysis of noisy spectrometric data, Savitzky and Golay developed a technique (Savitzky and Golay, 1964) which meets our requirement. **Savitzky-Golay (SG)** filtering can be thought of as a generalized weighted moving average. The weights are given as a polynomial of a certain degree. When SG is applied to a signal, it performs a polynomial least-squares fit within the filter window. This polynomial is designed to preserve higher moments within the data and reduce the bias introduced by the filter. This filter works especially well when the typical peaks of the signal are narrow. The heights and widths of the curves are generally preserved.

In the implementation of the 'Yaros' software, Christian Marquardt has designed a way to smooth and to thin data in a coherent manner using the SG algorithm. The SG filter (ng, nl, m) is used in the time domain to reduce the data and — with the same configuration — to sample every $(ng+nl+1)$ -th points. This is a very good way to achieve both smoothing

and thinning in a coherent manner.

Advantages

- The SG is a good way to achieve both thinning and smoothing in a coherent manner
- SG is very computationally efficient.

Disadvantages

- As for any sampling, the results are no longer defined on a fixed set of altitudes but on levels which depend on the profile.

An implementation of the SG filter can be found on the web at:
<http://www.library.cornell.edu/nr/bookcpdf/c14-8.pdf>.

An SG filter is implemented in IDL, as procedure *savgol.pro*, based on the Savitzky-Golay Smoothing Filters described in section 14.8 of "Numerical Recipes in C".

2.6 Conclusion

Although the most efficient and elegant way to smooth the data and thin it at the same time is the adoption of a SG filter with sampling, the output of this procedure is not defined on a fixed set of altitudes. Having a fixed set is a required feature, since it simplifies the intercomparison of profiles. ECMWF have stated a strong user requirement for such a feature of GRAS SAF NRT products in BUFR and prefer the set of impact height levels shown in Chapter B. This implies that the last step of the thinning procedure has to be an interpolation on fixed altitude (impact height) levels.

The preferred interpolation procedure would be a cubic spline. The spline does not smooth data and so it can be preceded by a SG smoothing filter which optionally can also produce an accessory uninterpolated output. Examples of the performance various thinning methods is shown in Figure 2.6

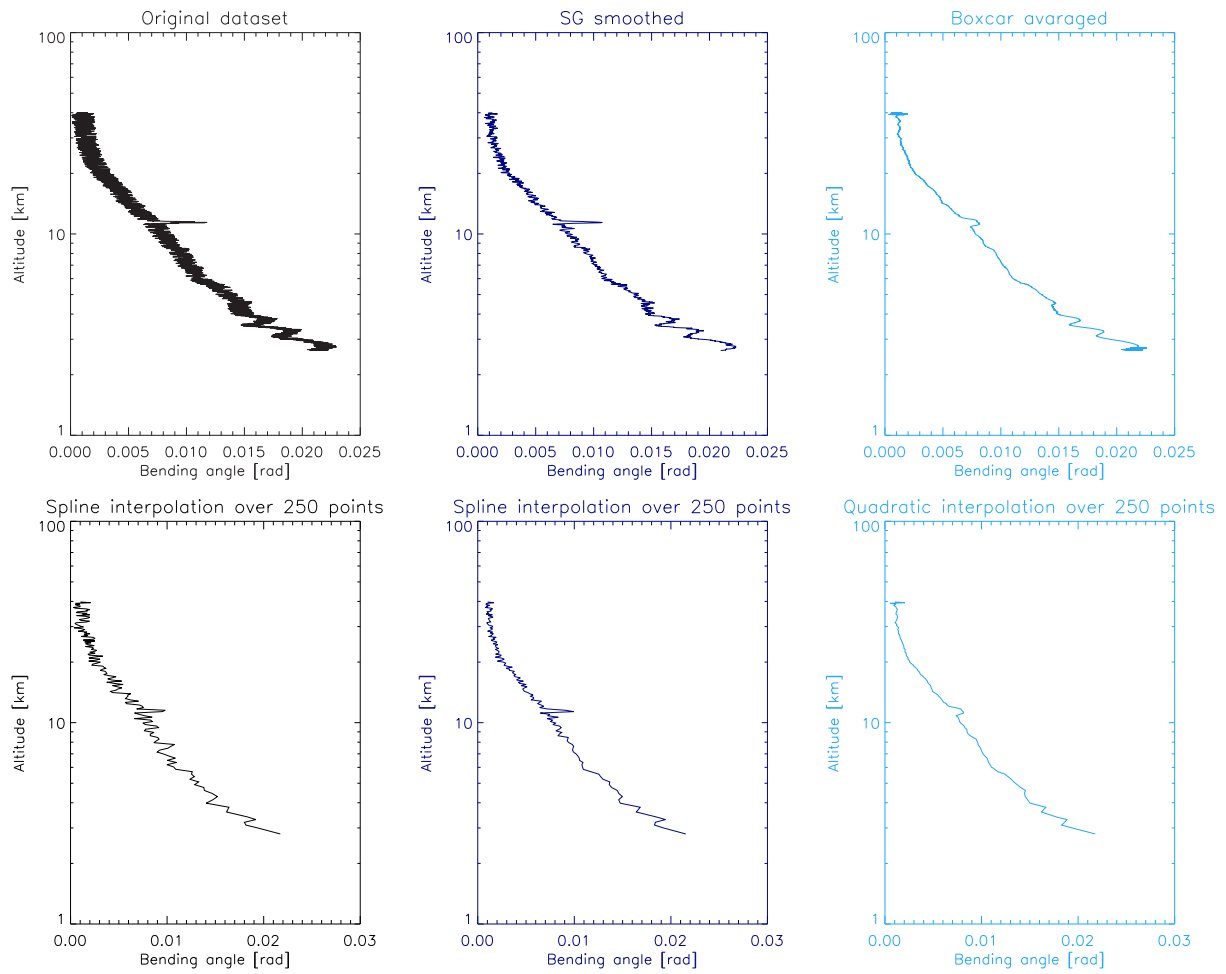


Figure 2.6: The effect of a quadratic spline interpolation for different thinning approaches

Bibliography

- Dyn, N., Floater, M., and Iske, A., *Univariate adaptive thinning. Mathematical Methods for Curves and Surfaces*, T. Lyche and L.L. Schumaker (eds.), Vanderbilt University Press, Nashville, 2000.
- Gorbunov, M., Gurvich, A., and Kornblueh, L., Comparative analysis of radioholographic methods of processing radio occultation data, *Radio Sci.*, 35, 1025–1034, 2000.
- Hinson, D. and Magalhaes, J., Equatorial waves in the stratosphere of Uranus, *Icarus*, 94, 64–91, 1991.
- Jensen, A., Lohmann, M., Benzon, H.-H., and Nielsen, A., Full Spectrum Inversion of radio occultation signals, *Radio Sci.*, 38, 1040, doi:10.1029/2002RS002763, 2003.
- Kursinski, E. R., Hajj, G. A., Schofield, J. T., Linfield, R. P., and Hardy, K. R., Observing earth's atmosphere with radio occultation measurements using the global positioning system, *J. Geophys. Res.*, 102D19, 23 429–23 466, 1997.
- Marquardt, C. and Healy, S., Measurement noise and stratospheric gravity wave characteristics obtained from GPS occultation data, *J. Met. Soc. Jap.*, 83, 2005.
- Rodgers, C., *Inverse methods for atmospheric sounding: Theory and practice*, World Scientific Publishing, Singapore, New Jersey, London, Hong Kong, 2000.
- Savitzky, A. and Golay, M., Smoothing and differentiation of data by simplified least square procedures, *Anal. Chem.*, 36, 1627, 1964.
- Sokolovskiy, S., Modeling and inverting radio occultation signals in the moist troposphere, *Radio Sci.*, 36, 441–458, 2001.

A Glossary

- Smoothing

An explicit attempt to reduce random noise and which inevitably introduces artifacts. It has no way to distinguish between real gradient in the dataset and noise. Generally smoothing also induces vertical correlations. If the behaviour of the smoothing procedure is well characterised it is possible to take account of it in any further processing. Smoothing does not necessarily imply thinning. It is a destructive operation since the data values are altered and once the data have been smoothed there is no way to come back to the original "noisy" profile.

- Thinning

A procedure which reduces the data volume by *generic* means, whilst trying to minimise loss of information. Thinning does not necessarily imply smoothing, and may be destructive or non-destructive.

- Sampling

A thinning procedure used to identify a sub-set of the original dataset by selecting some particular points. The sampling rate could be constant (every N -th point) or based on some statistical property of the dataset itself. Sampling is the basis of any "pure thinning" algorithm.

- Interpolation

A procedure which evaluates the values of the input (typically the measured parameter) on a specific set of output levels. If the number of output levels is much smaller than the number of input data points then the interpolation induces a thinning of the dataset. In radio occultation the bending angle profile is usually largely over-sampled, in terms of independent pieces of information available, so interpolation does not significantly reduce the information content. For the same reason it does not produce excessive smoothing. Interpolation is destructive as the original profile cannot be recovered. In the ROPP implementation the raw bending angles are interpolated onto a set of 247 impact heights for encoding into the BUFR format for dissemination over the GTS.

- Statistical optimisation

A variational retrieval of the bending angle profile which assumes an *a priori* knowledge derived from climatology Healy and Eyre (2003), of the bending angle profile. This is

a largely destructive procedure which alters the observed profile in order to do an optimal retrieval (in a statistical sense). Such pre-optimisation would be unnecessary if a realistic covariance matrix for the unoptimised profile were available (Hinson and Magalhaes, 1991).

Bibliography

Healy, S. and Eyre, J., A forward model for the assimilation of GPS radio occultation bending angle measurements, *Quart. J. Roy. Meteorol. Soc.*, submitted, 2003.

Hinson, D. and Magalhaes, J., Equatorial waves in the stratosphere of Uranus, *Icarus*, *94*, 64–91, 1991.

B Thinning Levels

The list of impact heights used for thinning the data with ROPP SG is shown below. This list has been defined by C. Marquardt, using the criteria of having a maximum of four points per Fresnel diameter. This set is endorsed by S. Healy as being the set of thinned bending angle fixed levels required by ECMWF for GRAS SAF BUFR products.

2094.240	2211.910	2330.080	2449.710	2569.340	2690.920
2812.500	2935.550	3058.110	3183.110	3308.110	3433.590
3560.550	3687.500	3815.920	3944.340	4073.240	4203.120
4333.980	4465.330	4597.170	4729.980	4862.790	4997.070
5131.350	5266.110	5401.860	5538.090	5675.290	5812.990
5951.170	6090.330	6229.980	6370.120	6510.740	6652.830
6794.920	6937.990	7082.520	7226.070	7371.580	7517.090
7664.060	7811.040	7958.980	8107.910	8257.320	8407.230
8559.080	8710.940	8863.280	9017.090	9171.390	9326.660
9483.400	9639.650	9797.360	9956.050	10116.70	10276.90
10438.50	10600.60	10764.20	10927.70	11093.30	11260.30
11426.80	11595.20	11765.10	11935.50	12106.90	12279.80
12453.60	12627.90	12804.20	12981.00	13159.20	13338.90
13519.50	13701.70	13884.30	14068.80	14254.90	14441.90
14629.40	14818.40	15010.30	15201.20	15395.50	15590.80
15786.60	15984.40	16182.60	16383.80	16585.90	16789.10
16994.60	17200.70	17408.20	17617.70	17828.60	18041.00
18255.40	18470.70	18689.00	18907.70	19128.90	19350.10
19574.20	19800.30	20027.30	20256.80	20486.80	20718.80
20953.10	21189.50	21426.30	21665.00	21905.80	22149.40
22393.60	22639.20	22887.70	23136.70	23388.70	23642.10
23897.00	24153.30	24412.10	24671.90	24933.10	25196.80
25462.40	25728.50	25997.60	26267.10	26539.10	26813.00
27088.40	27364.70	27643.10	27922.40	28203.60	28486.30
28770.00	29056.60	29342.80	29631.80	29921.40	30212.40
30503.90	30797.40	31092.80	31387.70	31684.60	31982.40
32281.20	32581.10	32881.30	33182.60	33484.90	33787.60
34091.30	34396.00	34700.70	35005.90	35311.00	35617.70
35923.30	36230.00	36537.60	36843.80	37150.90	37458.50
37765.10	38072.30	38379.40	38687.00	38994.60	39301.30
39608.40	39915.50	40223.10	40530.80	40836.90	41144.50
41452.10	41758.80	42065.90	42373.50	42681.20	42988.30
43295.40	43603.00	43910.60	44218.30	44524.40	44832.00
45139.60	45447.30	45754.40	46061.50	46369.10	46676.80
46983.90	47291.50	47599.10	47906.70	48214.40	48520.50
48828.10	49135.70	49443.40	49750.50	50058.10	50365.70
50673.30	50980.50	51288.10	51595.70	51903.30	52210.00
52517.60	52825.20	53132.80	53439.90	53747.60	54055.20
54362.80	54669.90	54977.50	55285.20	55592.80	55899.90
56207.50	56515.10	56822.80	57129.90	57437.50	57745.10
58052.70	58360.40	58667.50	58975.10	59282.70	59590.30
59896.50	/	/	/	/	/

GRAS SAF Reports

SAF/GRAS/METO/REP/GSR/001	Mono-dimensional thinning for GPS Radio Occultation
SAF/GRAS/METO/REP/GSR/002	Geodesy calculations in ROPP
SAF/GRAS/METO/REP/GSR/003	ROPP minimiser - minROPP
SAF/GRAS/METO/REP/GSR/004	Error function calculation in ROPP
SAF/GRAS/METO/REP/GSR/005	Refractivity calculations in ROPP
SAF/GRAS/METO/REP/GSR/006	Levenberg-Marquardt minimisation in ROPP
SAF/GRAS/METO/REP/GSR/007	Abel integral calculations in ROPP
SAF/GRAS/METO/REP/GSR/008	ROPP thinner algorithm
SAF/GRAS/METO/REP/GSR/009	Refractivity coefficients used in the assimilation of GPS radio occultation measurements

GRAS SAF Reports are accessible via the GRAS SAF website <http://www.grassaf.org>.