

Instruments and Methods

Singular spectrum analysis and envelope detection: methods of enhancing the utility of ground-penetrating radar data

John C. MOORE, Aslak GRINSTED

*Arctic Centre, University of Lapland, Box 122, FIN-96101 Rovaniemi, Finland
E-mail: jmoore@ulapland.fi*

ABSTRACT. We present a novel method of improving signal-to-noise ratio in radargrams. The method uses singular spectrum analysis (SSA) to separate each individual radar trace into orthogonal components. The components that explain most of the original trace variance contain mainly physically meaningful signal, while the components with little variance tend to be noise. Adding the largest-magnitude components together until the sum of components accounts for the variance above the noise level (typically 60–80%) of the original trace variance results in a much cleaner radargram with more easily seen internal features than in traditionally filtered data. The radargrams can be further enhanced by envelope-detecting the SSA-filtered data, as this measure of instantaneous energy minimizes the deleterious effects of innumerable phase changes at dielectric boundaries. Subsequent incoherent stacking results in far more structured radargrams than are achieved with traditionally processed radar data and amplitude stacking.

INTRODUCTION

Ground-penetrating radar (GPR) is a very widely used glaciological tool. Commonly GPRs are impulse radars that operate between a few MHz and 1 GHz (e.g. Sinisalo and others, 2003). Lower frequencies give greater penetration but lower resolution, while higher frequencies gain resolution at the expense of much less penetration (e.g. Watts and England, 1976). Frequently it is not simply the depth to bed that is of primary interest, but the internal reflections present in the ice (e.g. Miners and others, 1998; Eisen and others, 2003; Sinisalo and others, 2003). These reflections are usually much weaker than the bed reflector, and usually require more processing power and effort to recover. The vast amount of data usually collected during surveys is seldom exploited beyond the standard output available from fairly elementary qualitative processing and interpretation software such as Radan (Geophysical Survey Systems Inc.) or Haescan (Roadscanners Oy). In glaciological research there are frequently prohibitive logistical difficulties in returning to a particular glacier to collect additional data that may aid interpretation. Therefore it is worth going to more trouble than would be the case in a commercial survey to extract the maximum glaciological information from existing radar data. Typically, radar data processed with commercial packages, which are aimed more at the commercial user than at the researcher, are often manipulated in ways that are not easy to understand as the software source code is private. This severely limits the additional processing researchers can do on the output from the commercial software.

Here we present a simple method that is easily programmed (e.g. in Matlab[®]), and which can be used on data in a complementary way to commercial software, either on raw data or on output that has been manipulated by other software.

METHODS

Singular spectrum analysis (SSA) was developed as a means of analyzing time series (Allen and Smith, 1996; Ghil and others, 2002). The method is particularly suited to short, noisy time series such as those often found in palaeoclimate datasets. An individual radar trace spans the moment the radar transmitter antenna is triggered to the end of the time period that contains useful reflections. These radar traces are typically only about 1 μ s in length, yet they can be analyzed by any time-series method. The method we describe is not a spectral technique; it operates in the time domain, and as such is suited to time series containing quasi-periodic signals, rather than strict sinusoids. This is the situation with a radar trace produced by reflections of the transmitted wavelet from the virtually infinite number of subtle dielectric boundaries in a real glacier. The recorded trace contains a continuously varying power spectrum as the radar waveform is modified over its propagation path through the glacier by the constantly changing ice dielectric properties. This is the case in polar ice, despite density contrasts being rather smooth, as the dielectric constant of firn changes by a factor two in the upper 100 m of the ice sheet due to density increase alone. Thus, accurate modelling of radar traces in Antarctica requires a time-domain finite-difference solution (e.g. Miners and other, 1998; Eisen and others, 2003). In temperate or polythermal glaciers, where melt layers cause rapid changes in density and dielectric properties, the radar wavelet changes even more dramatically with depth. This means that simple frequency filtering is unlikely to be a good way of increasing the signal-to-noise ratio, and in our experience it seldom clarifies radargrams.

We decompose each of the radar trace time series using SSA (Allen and Smith, 1996) with a set of data-adaptive, orthogonal filters. SSA amplifies signal-to-noise ratio by separating the original time series into low-frequency trends




Fig. 1. SSA reconstructed components found by projecting the SSA filters found using the whole 2000 traces in Figure 4, on trace number 1, ordered by magnitude of variance accounted for in the radar trace. The noise floor we estimate as being after the first six components, which accounts for 63% of variance.

and narrowband quasi-periodic signals, with the rest (assumed to be noise) distributed among the filters. The number of orthogonal filters used is called the embedding dimension and is an important parameter in SSA. SSA operates by examining the dynamics of a complex system using the time-lagging method (e.g. Abarbanel, 1996). In this method the only information on the dynamics of the system comes from a time series of observations, in this case the voltages on the receiver antenna. A phase space that will approximate the dynamics of the system is constructed from the coordinates of the original time series and delayed copies of it. Each copy is delayed by a different integer multiple of time known as the lag. The lag is often, as it is here, the time interval between successive samples in the time series. The total number of time-series copies defines the embedding dimension. The smaller the embedding space, the shorter the length of the window over which the resolved components are calculated, and the less resolved is each component. The longer the window, the greater the frequency resolution of each component, but the greater the chance that noise is mistaken for signal and a greater proportion of the time series is affected by the data boundaries. The embedding dimension also determines the low-frequency limit below which oscillations will appear as trends (Moore and others, 2005). In typical time-series analysis we would calculate the statistical significance of individual components by testing the variance explained by the component against suitable noise models using Monte Carlo methods. The Monte Carlo methods are much the slowest computational part of the procedure; however, in the case of radar data we do not need to do this as we have many time series available and the significance of signals can be directly assessed by the observer.

The set of incrementally time-lagged radar traces make up the lagged-data matrix, representing the dynamical phase

space. The basic idea behind SSA is to linearly transform the highly correlated set of vectors in the lagged-data matrix into a set of vectors that are linearly independent. The dependence of the lagged data can be expressed in terms of the covariance matrix (C) of the lagged-data matrix. The eigenvectors (p_k , $k = 1 \dots M$) of C are the principal directions in the data matrix and are often called empirical orthogonal functions (EOFs), where M is the embedding dimension. Principal components (PCs) are obtained by projecting the lagged-data matrix onto the eigenvectors. The eigenvalues (V_k) of C tell us how much variance is captured by each PC and the eigenvectors and eigenvalues are sorted by importance so that V_1 is largest. The PCs that capture little of the variance (low eigenvalues, V_k) can be considered noise. The original time series can be reconstructed from the full set of PCs and EOFs. This can be done so that the reconstructed series is a sum of a set of reconstructed components (RCs) where each RC is only derived from one PC and its associated EOF. Hence, it is possible to make a partial reconstruction of the time series where components that capture little variance and are thus thought to be noise can be excluded. We can write the reconstructed components as

$$R_k(t) = \frac{\rho_k(t) * \rho_k(-t) * X(t)}{M},$$

where $*$ indicates convolution, t is time and X is the time series. For radar filtering, the eigenvectors can be derived from the lag covariance matrix either of the full radargram or on a trace-by-trace basis. Figure 1 shows the RCs found for a single radar trace, and the associated eigenvectors derived from the covariance matrix of 2000 radar traces.

A wise choice of embedding dimension can be made with a priori insight. However, our experience with radar data suggests this choice is not particularly critical as we are mainly interested in the noise reduction properties of SSA as opposed to examining the particular SSA RCs that may correspond with physically meaningful oscillations or trends in the radar data. As larger embedding dimensions carry a computational penalty and are affected more by the data boundaries, we use rather small values of M : typically a value of 20 or 30 is sufficient to perform adequately. We recommend that a wide range of M is experimented with to examine possible unwanted artefacts.

For radar data, we simply sum the reconstructed components that account for the fraction of the original time-series variance that is above the noise floor of the data. The noise floor can be found by plotting the magnitude of V_k as a function of their index k when ordered by magnitude (Fig. 1), or equivalently by considering the contribution to the variance of the original radar trace of the RCs. Often there will be rapid fall-off of V_k with k until a fairly stable level is reached. This stable level represents the noise floor, and the components at this level are then representing mainly noise. The variance accounted for by the number of components above the noise floor is then used as the number of RCs to be used in the SSA filtering. Typically, in our experience, the useful signal is 60–80% of the total variance in the radar trace. Figure 1 exhibits two rapid drops in variance (after the first two components and after the third and fourth components), then there is a fairly gradual drop-off in variance accounted for by the RCs rather than a low-level plateau. This example presents a more

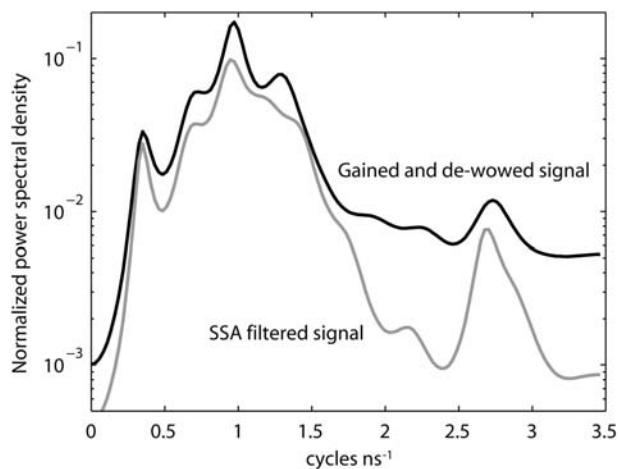


Fig. 2. The normalized power spectral density found using the maximum entropy method of order 20, for a block of 50 traces of radar data used to make one stacked trace in Figure 5. Note the peak in power at around 2.6 cycles ns^{-1} (2600 MHz) that is kept as signal by SSA that would very likely be rejected by bandpass-filtering the 800 MHz centre frequency radar data.

difficult situation than arises when the noise floor is very clear, but there is still only a slow drop-off in variance accounted for beyond the sixth component. Here, we somewhat arbitrarily select six components as the cut-off of signal from noise, equivalent to accounting for a variance of 63%.

Removal of the non-linear trend component in SSA (e.g. Moore and others, 2005) corresponds with the 'd.c. removal' or 'de-wowing' processes usually done in radar processing, whereby the shift in reference voltage level and very long-period oscillations in each radar trace are removed. The SSA approach has the advantage of being a purely objective method of selecting parts of the received antenna voltage power spectrum that correspond with meaningful signals. A common method of attempting a similar process is to use bandpass filtering. However, particular SSA filters often correspond with rather narrowband quasi-period oscillations, and they can also come from anywhere in the full-frequency spectrum of the received radar trace rather than representing high- and low-pass limits. The difference between SSA and bandpass filtering can be seen by comparing the power spectrum of the amplified radar received data with the power spectrum of the SSA components that together comprise the reconstructed signal (Fig. 2).

ENVELOPE DETECTION STACKING

One common method of improving signal-to-noise ratio in radar is stacking. This method relies on coherent signals adding, while noise, being incoherent, rises only as the square root of the stacking number. This should occur where radar data are collected with the antennas stationary. Usually, however, radar data are collected while the antennas are continuously being moved over the snow surface. Thus stacking is done by averaging the radar traces collected over an appreciable distance of horizontal travel. When the radar frequencies used are high, and the dielectric properties of the firn and ice vary rapidly laterally, the assumption of coherence may fail. Since the bandwidth of




Fig. 3. The steps in the processing of radar trace No. 1000 in Figure 4. (a) The originally recorded radar trace; (b) after removal of the mean background trace; (c) simple linear gaining; (d) after removal of the noise SSA components; and (e) after envelope detection. Signal scaling is arbitrary.

impulse GPR is high, changes in dielectric layers, even of a few centimetres over distances between radar trace collection, (commonly some decimetres) can cause undesirable phase mismatching in stacking. In our experience, this often occurs, and stacking usually fails to bring significant improvement in signal-to-noise ratio.

Signal envelope detection (Nye and Berry, 1974) is a way of processing the data such that they can be stacked in an incoherent way, but which does significantly improve the signal-to-noise ratio. This method essentially squares and sums the real measured antenna voltage, and its quadrature, that is equivalent to the magnitude of the Hilbert transformation of the radar trace. In practice, this can be simply and very efficiently implemented by summing the squared radar trace data and their squared numerical differences, since the differential or integral of a sinusoid is just a phase shift of $\pi/2$. In some commercial processing packages, the magnitude of the Hilbert transformation is directly implemented. The envelope calculated represents the energy around the particular data point in the radar trace, so the envelope is not affected by the phase changes caused by the many dielectric discontinuities in the firn. Therefore small-scale lateral inhomogeneity that can undermine coherent stacking does not adversely affect the envelope-detected stack, and hence its usefulness for GPR data.

PRACTICAL EXAMPLES OF THE METHOD

Figure 3 shows all the steps in processing a radar trace that we use in the example radargrams discussed next.

To illustrate the advantages of the method we describe over normal processing, we show some radar data collected in Antarctica during December 1999 (Sinisalo and others, 2003). The first example is of a radar survey on Plogbreen,




Fig. 4. (a) A profile of 800 MHz data from Antarctica (Sinisalo and others, 2003), that has had the mean background trace removed and linear gaining applied. (b, c) The same image after SSA-processing (b) and after envelope-detecting (c) the radar traces.

about 6 km from the Finnish research station Aboa ($73^{\circ}03' S$, $13^{\circ}25' W$). The radar used was a Malå Geoscience Ramac system with an 800 MHz centre frequency shielded-antenna system. The radar control unit, computer and a roving global positioning system (GPS) receiver were mounted on the snowmobile. The antenna system was pulled about 2 m behind the snowmobile, using rigid struts to keep it at a fixed distance. No trace stacking was done and data were collected on a laptop computer. Each trace contains 1024 samples spanning a time window of 197 ns. Traces were collected at 0.25 s intervals, which correspond to horizontal distances of about 0.8 m at the snowmobile travelling speed. We reproduced, using Matlab[®], the processing done by Sinisalo and others (2003). Background noise was removed and vertical bandpass filtering was performed. Figure 4 shows an example containing 2000 traces from the radar line, a distance of about 1600 m.

In our new method we do the same initial processing steps on each trace (see Fig. 3), but instead of the frequency filtering we perform SSA. As discussed above, we select an embedding dimension that is relatively short; here we use a value of 20. Figures 1 and 3 show the results of applying SSA to a single trace. The larger-magnitude reconstructed components in Figure 1 have clear changes in amplitude, as may be expected at an internal reflection horizon, but many of the lesser-magnitude components also possess that feature. We think that much more could be achieved by analyzing these components in more detail, as has been done with SSA quasi-periodic components in palaeoclimate time series, which reveal rich structure in the underlying climate system (e.g. Jevrejeva and others, 2004).

Figure 1 suggests that about 60–70% of the variance is above the noise floor and that the six largest reconstructed




Fig. 5. (a) Radar data processed by background removal, d.c. removal, gaining, de-wowing and high-pass filtering, and stacking 50 times. (b) The same data processed by background removal, gaining, SSA filtering, envelope-detecting and stacking 50 times.

components amount to 63% of the original trace variance. We therefore extract the six leading RCs from all the 2000 traces in the section shown in Figure 4, to give the SSA-processed image in Figure 4b. There is clearly an enhancement in the SSA-processed image in Figure 4b relative to the normally processed data of Figure 4a. It is clear that the radar data contain eight to ten gently convex-up internal reflection bands, with hints of more detailed structure in places. In particular, there are reflecting layers in the deeper half-profile visible between traces 1500 and 2000 that are completely lost in the normally processed data. The layering is even more obvious in the envelope-detected image in Figure 4c.

It is of interest to consider the cause of the layer reflections. Sinisalo and others (2003) show that accumulation rates in the area are about $300 \text{ mm w.e. a}^{-1}$, and this would be about 500 mm of snow and firn in the upper layers. Given a radar wave velocity of $200 \text{ m } \mu\text{s}^{-1}$, 500 mm thick layers are about 5 ns of two-way travel time. Thus the layering is not simple annual layers, but more likely variation in densities due to the occasional melt events in the area.

The second example we consider examines the advantages of stacking GPR data. Again we consider 800 MHz Ramac radar data collected using the same set-up as described earlier. Each trace consisted of 1024 samples that were collected in a time window of 148 ns, with traces collected at 0.1 s intervals, which corresponds with three traces per metre. The data come from an area ($74^{\circ}34' S$, $10^{\circ}43' W$) on Ritscherflya, about 15 km from the Swedish Svea station and near to the site of a Dutch automatic weather station. In total, 23 384 traces were collected, and we process the data as described earlier, but with 50 times stacking. SSA processing showed that about 80% of data were above the noise floor. The reconstructed components accounting for 80% of variance were summed and envelope-filtered and then stacked 50 times. Figure 5 shows the results: no layering can be seen in the normally processed and stacked radargram (Fig. 5a), while layers are quite easily seen in the SSA- and envelope-processed data (Fig. 5b). In this example, the envelope detection part of the processing is responsible for most of the improvement,

mainly because there was relatively little noise present in the radar data for SSA filtering to remove. Stacking was needed in this example to reveal layering because of the very small density differences present in the snow at this site. No seasonal melting occurs and wind speeds are high, ensuring considerable surface sastrugi and disturbance to annual layering. The layers are about 5 ns apart in Figure 5, suggesting that they, in contrast with the layers in Figure 4, probably represent annual stratification.

CONCLUSIONS

We present a novel method of processing radar data that should be especially well suited to impulse GPR systems. The SSA method is superior to common bandpass filtering approaches to improving signal-to-noise ratio in several ways: (1) an objective method of separating signal and noise components of the radar trace comes from analysis of the variance accounted for by the orthogonal reconstructed components; (2) SSA non-linear trend terms correspond with the separate 'd.c. level' and 'de-wowing' corrections commonly done to radar data; (3) the full power spectrum of the radar trace is available for analysis rather than arbitrarily setting bandpass filtering limits. The individual SSA reconstructed components also merit examination separately as they contain information on aspects of the complex system comprising the electromagnetic wave interaction with the dielectric medium of sub-parallel firn and ice layers.

We show that the GPR traces that have been 'envelope-detected' can be stacked much more successfully than can be done simply by amplitude stacking as a means of improving signal-to-noise ratio. This is because in high-bandwidth GPR data small lateral changes in the dielectric properties of firn layers cause phase shifts large enough to destroy coherent amplitude stacking in many practical surveys. Envelope detection is also unaffected by electromagnetic phase changes on reflection, as it is a measure of the instantaneous energy at particular times in the radar trace.

ACKNOWLEDGEMENTS

The field logistics in 1999–2002 were provided by the Finnish Antarctic Research Programme (FINNARP) 1999–2002, and we thank A. Sinisalo for help with radar data collection. The work was funded by the Finnish Academy and the Thule Institute. B. Hubbard and R. Pettersson provided helpful comments on the manuscript.

REFERENCES

- Abarbanel, H.D.I. 1996. *Analysis of observed chaotic data*. New York, Springer.
- Allen, M.R. and L.A. Smith. 1996. Monte Carlo SSA: detecting irregular oscillations in the presence of colored noise. *J. Climate*, **9**(12), 3373–3404.
- Eisen, O., F. Wilhelms, U. Nixdorf and H. Miller. 2003. Revealing the nature of radar reflections in ice: DEP-based FDTD forward modeling. *Geophys. Res. Lett.*, **30**(5), 1218–1221.
- Ghil, M. and 10 others. 2002. Advanced spectral methods for climatic time series. *Rev. Geophys.*, **40**(1), 1003. (10.1029/2000RG000092.)
- Jevrejeva, S., J.C. Moore and A. Grinsted. 2004. Oceanic and atmospheric transport of multiyear El Niño–Southern Oscillation (ENSO) signatures to the polar regions. *Geophys. Res. Lett.*, **31**(24), L24210. (10.1029/2004GL020871.)
- Miners, W.D., N. Blindow and S. Gerland. 1998. Modelling a GPR record by using ice core data to produce a synthetic radiogram. *In Proceedings of the Seventh International Conference on Ground-Penetrating Radar, May 27–30 1998, Lawrence, Kansas, USA*. Lawrence, KS, Radar Systems and Remote Sensing Laboratory, 437–442.
- Moore, J.C., A. Grinsted and S. Jevrejeva. 2005. New tools for analyzing time series relationships and trends. *EOS Trans. AGU*, **86**(24), 226, 232.
- Nye, J.F. and M. Berry. 1974. Dislocations in wave trains. *Proc. R. Soc. London A*, **336**, 165–190.
- Sinisalo, A., A. Grinsted, J.C. Moore, E. Kärkäs and R. Pettersson. 2003. Snow-accumulation studies in Antarctica with ground-penetrating radar using 50, 100 and 800 MHz antenna frequencies. *Ann. Glaciol.*, **37**, 194–198.
- Watts, R.D. and A.W. England. 1976. Radio-echo sounding of temperate glaciers: ice properties and sounder design criteria. *J. Glaciol.*, **17**(75), 39–48.

MS received 11 February 2005 and accepted in revised form 4 November 2005