UNIVERSITY OF SOUTHAMPTON

Department of Aeronautics and Astronautics

School of Engineering Sciences

THE DEVELOPMENT OF A SOLAR PROTON EVENT PREDICTION MODEL

by

Gareth James Patrick

Thesis for the degree of Doctor of Philosophy

May 2003

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS

SCHOOL OF ENGINEERING SCIENCES

Doctor of Philosophy

THE DEVELOPMENT OF A SOLAR PROTON EVENT PREDICTION MODEL

by

Gareth James Patrick

Solar Proton Events (SPEs) pose a significant radiation hazard to spacecraft within and beyond the magnetosphere, but currently there is no capability to predict these events more than 24 hours before they occur. This thesis develops a classification approach for the prediction of SPEs with a 48-hour lead time, and addresses the fact that very little work has been done on examining SPE forecasting methods with longer lead times than current flare-association techniques allow. Development of the technique has been based on a uniform dataset that covers 3 solar cycles and more than 30 decades of continuous spacecraft observations, and has used solar x-ray fluxes and solar radio fluxes as predictor variables.

By comparing times of SPE occurrence to times at which the solar proton flux was at a background level it has been shown that SPEs are associated with increased levels of solar x-ray flux and solar radio flux, and that these increases are, on average, significant up to 5 days prior to SPE occurrence. Using these variables as inputs neural models have generated 65% success rates for SPE prediction with a 48-hour lead time, extending the lead time of existing models by a day or more. A neural model has been coded to operate in real-time and represents the only autonomous SPE forecast model with a 48-hour lead time that does not require human supervision. Assessing the model over a 12-month operational period showed it to have superior SPE detection capability to the current 2-day forecast operated by the Space Environment Centre.

Success of the classification technique was limited by the fact that solar x-ray flares were found to exhibit similar precursors to SPEs, although this meant that the model could in fact be used to forecast flares to a greater success than SPEs. Additional findings showed that the correction of radio flux observations for centre-to-limb dependence may offer the potential for more accurate forecasting ability on a timescale of days.

TABLE OF CONTENTS

APPENDICES

REFERENCES

BIBLIOGRAPHY

LIST OF FIGURES

LIST OF TABLES

DECLARATION OF AUTHORSHIP

I,..........Gareth James Patrick..

Declare that the thesis entitled:

The Development of a Solar Proton Event Prediction Model

and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at The University of Southampton;
- where any part of this thesis has been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- Parts of this work have been published as:

Patrick G., Gabriel S., "Neural Network Prediction of Solar Proton Events with Long Lead Times". *Euroconference: Solar Cycle and Space Weather*, Vico Equense, Italy, 24th - 29th September 2001.

Patrick G., Gabriel S., "Neural Network Prediction of Solar Proton Events". *Space Radiation Environment Workshop*, DERA Farnborough, England, 1st – 5th November 1999.

Patrick G., Gabriel S., "Prediction of Solar Proton Events using Neural Networks". *SHINE Conference*, Boulder, Colorado, June 1999.

Signed..

Date..........March 2004..

Acknowledgements

For their help in providing use of RBF modelling software and for their invaluable advice on JAVA programming, I would like to thank David Rodgers and Simon Clucas of QinetiQ. I am also indebted to Dr. Stephen Gabriel for his guidance throughout the past few years and his valuable suggestions in preparing this thesis. I would also like to say a special thank-you to my partner for her patience and ongoing encouragement. Finally, I would like to express special appreciation to my fellow researchers from the Astronautics Research Group at Southampton University for their continued support and motivation.

ABBREVIATIONS

Å	Angstrom (10^{-10} m)
ARIMA	Auto-Regressive Integrated Moving Average
ESA	European Space Agency
EVA	Extra-Vehicular Activity
GOES	Geostationary Observational Earth Satellite
IMP	Interplanetary Monitoring Platform
MINITAB	Commercial Statistics Software Package
MLP	Multi-Layer Perceptron
PFU	Proton Flux Unit (protons.cm^{-2}.s^{-1}.sr^{-1})
POPE	Predictor Of Proton Events
RBF	Radial Basis Function
SEC	Space Environment Centre
SEM	Space Environment Monitor
SPE	Solar Proton Event
TSS	Training Set Size
XL	Long X-Rays (1-8Å)
XS	Short X-Rays (0.4-4Å)

1. INTRODUCTION

Solar Proton Events (SPEs) are sporadic solar emissions of high energy particles and pose a significant hazard to both earth orbiting and interplanetary spacecraft. The occurrence of an SPE can cause proton fluxes in geostationary orbit to rise rapidly by several orders of magnitude and remain elevated for several days, causing effects that range from increased errors in spacecraft memory to permanent latch-ups in electronic instruments. The high energy protons that constitute SPEs are also a large potential risk to future manned missions involving Lunar or Martian exploration, and their occurrence is also thought to contribute to radiation doses at aircraft altitudes.

Whilst relatively accurate models exist to estimate radiation doses in space over timescales of years there is little provision for the accurate real-time prediction of SPEs. Currently, only two real-time SPE forecast models are in use, operated by the Space Environment Centre (SEC), and their usefulness is limited in that both of these models require an x-ray flare to occur before an SPE prediction can be made. This fundamentally limits the lead time of SPE forecasts to several hours, and as a result their occurrence is often without adequate warning, and their impact to operations is high.

This thesis is concerned with the development of a new SPE prediction model which aims to make SPE predictions with longer lead times than current methods by using inputs that are not from discrete x-ray flares. The work addresses the fact that the prediction of SPEs without using x-ray flares has never been attempted, hence it is currently unknown whether or not SPE-precursors other than flares exist. The advantage in developing such an approach is that lead times are no longer limited to the physical time difference between x-rays and particles arriving at earth, (which is typically less than 12-hours), allowing more time for appropriate measures to be taken. The typical role of an SPE forecast service is to alert spacecraft operators to a likelihood of satellite anomalies, but if forecasts can be made with a suitable accuracy and lead time it may eventually be feasible to delay spacecraft launches or EVAs, and could serve as a warning to manned excursions on the lunar or Martian surfaces that may be several days travel from suitable radiation shelters.

The following work begins by looking at our current understanding of SPEs, and examines the models that are currently in use by the SEC to predict their occurrence. By evaluating the shortcomings of these existing real-time forecast models a blue print for a

new model is drawn, and then a number of different approaches are discussed as to how the new model might be implemented. This concludes with a preference towards empirical rather than physical modelling and the thesis goes on to develop both a time series forecasting approach and a classification approach involving neural networks.

It is concluded after applying autoregressive and moving average (ARIMA) techniques to a proton flux time-series that SPEs are better represented as discrete occurrences, which allow the proton flux to be represented as a binary time series. A classification approach has then been adopted in order to predict the time series by using solar x-rays and solar radio flux as predictors variables. This leads to the optimisation of an MLP configuration which uses the ratio of GOES solar x-ray fluxes as inputs to generate SPE forecasts with a 48-hour lead time. The model is subsequently coded to ESA software standards to operate in real-time and has been assessed over a 12-month period. It represents the only autonomous prediction model that currently exists with a 48-hour lead time.

The development of the classification model utilises a dataset spanning approximately 3 decades, and in addition to allowing the creation of a model, has also enabled the behaviour of solar x-ray flux and solar radio flux to be examined prior to the occurrence of over 100 SPEs. This has resulted in the tentative identification of possible longer term precursors to SPEs in solar x-ray and radio flux that have not been previously reported. Solar x-ray fluxes are shown on average to be higher prior to SPEs, and SPEs are shown to coincide with peaks in a 27-day periodicity in the 2800Mhz solar radio flux. It is also suggested that spatially resolved measurements of solar radio flux, which can be corrected for observer location, may reveal a more accurate correlation between radio flux peaks and SPE occurrence.

The thesis concludes by comparing the performance of the neural classification model with that of the SEC 2-day forecast over the same operational period, where it is shown that the neural model has far superior SPE prediction capability. The performance of the classification model is found to be limited by the fact that the solar x-ray flux and solar radio flux prior to flares with no SPE association exhibit similar behaviour as at prior to SPE occurrence. Findings from throughout the work and an examination of the classification models' behaviour are used to provide comments concerning directions of future research in this area, and include analysis of the model as an x-ray flare forecaster.

2. SOLAR PROTON EVENTS

This chapter reviews the current understanding of SPEs by examining their physical nature and their origin. The problems posed by SPEs are discussed, as are their known precursors and associations, and existing models for SPE forecasting and solar proton fluence estimation are examined in detail. The chapter ends by summarising the limitations of current real-time SPE forecasting tools, and provides evidence to suggest that alternative prediction approaches could be developed.

2.1 Solar Proton Monitoring

Solar Proton Events (SPEs) consist of atomic particles, predominantly protons, that are accelerated in the solar atmosphere to extremely high energies up to GeV ranges [*Shea and Smart*, 1994]. Solar protons have been monitored indirectly via muon and neutron monitors since circa 1930, but such techniques were only able to detect very high energy protons (500MeV and above) via secondary particles seen as increases in cosmic ray intensity. Atmospheric balloons and sounding rockets enabled more sensitive measurements to be taken, but since the mid 1960s (corresponding to solar cycle 20) solar particle fluxes have been monitored at energies ranging from 1 to 500 MeV by space borne instruments on board the Interplanetary Monitoring Platform (IMP) spacecraft and Geosynchronous Observational Environment Satellites (GOES) [*Shea and Smart*, 1995b].

The IMP and GOES platforms have provided valuable, uniform datasets over a continuous time span of decades (GOES platforms have been operational since 1974, and IMP platforms since 1963) and they have provided the basis for many solar proton studies over the past 30 years. The GOES satellites are used continually by the Space Environment Centre (SEC) to monitor solar x-rays and the proton flux environment in geostationary orbit, and it is often these measurements that first indicate that a solar proton event is in progress.

SPEs are typified by a rapid rise in the geostationary proton flux, usually by several orders of magnitude over a few hours, followed by a slow decay to background levels over a period of a few days. Das et al. studied proton events between 1955 and 1985 and reported rise times of up to 55 hours, but usually less than 20 hours, and durations of between 10 and 120 hours [*Das, Chakraborty, and Dasgupta*, 1987]. Events later in history however show that the large SPE of October 1989 caused the geostationary >10MeV proton flux to remain elevated for over 30 days.

There is no universally agreed definition as to what constitutes an SPE but the SEC defines one as occurring if the >10MeV proton flux (as measured by GOES satellites) is greater than 1.0 pfu, and remains above this level for 3 consecutive 5-minute observations. This definition appears to stem from user group requirements in that it is protons of energies greater than 10MeV that pose hazards to satellite components. The peak fluxes and fluences of SPEs vary over several orders of magnitude: on rare occasions >10MeV peak fluxes reach 10^4pfu, and >10 MeV event fluences of 10^{11} protons/cm^2 were recorded for the August 1972 event [*Feynman, Armstrong, Daogibner, and Silverman*, 1990b].

Shea and Smart performed a study based on highly energetic events (those detectable at ground level) between 1955 and 1986 but found no pattern in their occurrence other than a general association with solar activity in that there were more solar proton events during solar activity maximum than minimum [*Smart and Shea*, 1989]. Calendar years were found to contain between 1 and 16 SPEs depending on their position in a solar cycle. It was also noted that SPEs tended to occur in episodes of activity, with one active region producing a series of major flares with associated proton emissions as it passed across the central meridian.

2.2 The Effects of Solar Protons

The effects of solar activity on earth based systems were first recorded in 1860 when, due to induced currents from a geomagnetic storm, telegraph lines between Boston and Portland could be used without need for batteries [*Stewart*, 1861]. The increased reliance on electrical and space based technology means that today's systems are even more vulnerable to solar activity, one famous instance being the destruction of electrical transformers in Canada due to Ground Induced Currents (GICs) in 1989. Over the past 20 years the term 'Space Weather' has been coined to define the behaviour of Sun-Earth interactions, of which SPEs are just one facet of solar emissions, along with CMEs, radio and x-ray flares, and interplanetary shocks, all of which can be geo-effective. However, whilst SPEs are sometimes associated with geomagnetic storms, their largest impact is on the near earth space environment and the systems which operate within it.

The arrival of energetic particles at earth can cause the >10MeV proton flux at geostationary orbit to increase by several orders of magnitude and reach fluxes of 10^4pfu. Typical consequences experienced by spacecraft are a rapid and permanent degradation in solar panel performance, and an increase in the number of Single Event Upsets (SEUs) [*Shea, Smart, Allen, and Wilkinson*, 1992], [*Vampola, Lauriente, Wilkinson, Allen, and Albin*, 1994], [*Stassinopoulos, Brucker, Adolphsen, and Barth*, 1996]. SEUs are caused by single heavy ions penetrating sensitive areas of circuitry and can cause sufficient electrical charge to change the logic state of the device. SEUs often appear as memory errors and are reversible by rebooting onboard software, but if they occur in critical areas can result in permanent latch-ups, rendering instruments or sub-systems unusable [*Tranquille*, 1994].

Wilkinson reports that a family of geostationary spacecraft experienced 1 SEU per hour in relation to September 1989 solar proton emissions and in October of the same year GOES spacecraft suffered star sensor outage and major loss of solar panel output due to the direct result of another large SPE [*Wilkinson*, 1994]. The TDRS-1 spacecraft also exhibited significant increases in SEU rates relating to proton events during 1989, 1991 and 1992 [*Shea and Smart*, 1998], [*Wilkinson, Daughtridge, Stone, Sauer, and Darling*, 1991].

SPEs also pose a serious hazard to manned missions. It is calculated that the SPE of October 1989 would have almost tripled the current annual radiation allowance for NASA astronauts behind 2g/cm^2 of shielding [*Heckman, Kunches, and Allen*, 1991]. Heckman et al. have examined scenarios for possible future missions in which astronauts may be working on the surface of the moon or Mars, and suggest the requirement for a SPE prediction with a lead time suitable to allow a radiation shelter to be reached [*Heckman, Wagner, Hirman, and Hunches*,]. The PROTONS prediction model for SPE occurrence (operated by the Space Environment Centre) was developed prior to the Apollo lunar missions for the specific purpose of providing a hazard warning for lunar astronauts.

Very recently, the impact of SPEs and solar radiation on airline flights has become an issue with aircraft operators and radiological protection agencies. Research by O'Sullivan et al. addresses the fact that biological damage from neutrons may have been underestimated and states that the radiation doses received from flying at higher subsonic altitudes are unknown [*O'sullivan, Zhou, and Flood*, 2001], [*O'sullivan, Zhou,* *Heinrich, Roesler, Donnelly, Keegan, Flood, and Tommasino*, 1999]. A current study is flying high energy radiation monitors on selected commercial flights in order to quantify the radiation field, but the occurrence of a SPE is known to enhance the ambient particle conditions [*O'sullivan and Zhou*, 1999].

2.3 Solar Proton Modelling

Of particular importance at the beginning of the space era was the need to try and estimate the radiation dose that spacecraft would experience during their lifetime, and the availability of proton data from IMP satellites provided an opportunity for quantitative analysis of SPE fluence.

King was the first to use early satellite data to study SPE fluences, and using data pertaining to solar cycle 20 (1964-1976) he found that the radiation dose at 1AU was dictated entirely by contributions from SPEs, or, more specifically, the occurrence of a single, large fluence SPE (the August 1972 event). King went on to develop a model to predict the proton fluence for cycle 21, basing it on the events that had occurred during cycle 20 and the predicted sunspot maximum for cycle 21 [*King*, 1974]. Data pertaining to cycles 19 and 20 seemed to indicate that proton fluence was directly related to the maximum sunspot number of a cycle.

Time would show however that King had based his model on too small a dataset, and this is an inherent problem when attempting to draw conclusions from solar proton studies. Their low occurrence frequency and the further division of SPEs into their respective solar cycles means that decades must pass before reliable statements about their distributions can be made.

Following Kings' attempt at trying to predict proton fluences, later work by Feynman et al studied the distribution of fluences from SPEs occurring between 1956 and 1986, incorporating solar cycles 19, 20 and 21 [*Feynman, Armstrong, Daogibner, and Silverman*, 1990b]. In studying the distribution of SPE fluences Feynman made no distinction between solar cycles, treating the dataset as one population, with the result that the SPE fluences formed a continuous distribution. Exceptionally large fluence SPEs, previously thought by King to be anomalous events, were now shown to belong to the same distribution as smaller fluence SPEs. Feynman found that the distribution of event fluences could be approximated at higher fluences by a power law. Small (low fluence) events were found to occur far more frequently than high fluence events and prevented an accurate power law fit of the entire distribution, although the fitting functions (power laws) were biased to model the large fluence event probabilities as it is these events that dictate the proton fluence over a suitably long time span.

Feynman's study also revealed a pattern in that the significant majority of solar proton fluence was produced during a period spanning from 3 years before to 4 years after the time of solar sunspot maximum, defined to the nearest month. This finding is supported by Shea and Smart, who note that the majority of SPEs occur through the second through eighth years after sunspot minimum [*Smart and Shea*, 1989]. Feynman called this 7-year period the active years of a solar cycle and concluded that when estimating the proton fluence over a given time it was only necessary to consider the active years: the remaining years could be ignored as they gave only a small contribution to proton fluence in relation to active years. Using Monte Carlo methods Feynman generated curves based on the SPE distribution that showed the probability of exceeding a given proton fluence over different time spans (mission lengths). Feynman revises the work by including data up to 1991 and calculates probability curves for different energies, formalising the model as the JPL-91 Interplanetary Proton Fluence Model [*Feynman, Spitale, Wang, and Gabriel*, 1993], which is widely accepted as the industry standard for proton fluence estimation.

It is worth noting that the Feynman model was created in response to a need to be able to predict the radiation dose during a spacecraft lifetime. However, the notion of 7-active years is not realistic, and if a mission lasting 4 years happens to fall in the 4 non-active years between solar cycles it will have zero predicted radiation dose. Clearly, the Feynman model is only valid if the majority of a spacecraft mission falls into the seven active years of a cycle. In addition a study by Stassinopoulos et al. since the JPL-91 model has shown that 'non solar active years' can still contain significant SPEs, and in a few instances have also had higher annual proton fluences than supposed active years [*Stassinopoulos, Brucker, Nakamura, Stauffer, Gee, and Barth*, 1996]. This again highlights the fact that the concept of active years is really only valid as a probabilistic model over long time spans. It also shows that SPEs do occur outside 'active years', with the implication that a discrete SPE forecasting approach cannot limit itself to just the active years of a cycle.

More recent work has revised the proton fluence power laws [*Gabriel and Feynman*, 1996], [*Gabriel, Feynman, and Spitale G.*, 1996], [*Gabriel, Feynman, et al.* 1996] and others have used extreme value theory and Poisson distributions to generate more accurate and more simplistic probabilistic models for total dose prediction [*Xapsos, Summers, and Burke*, 1998], [*Xapsos, Summers, Shapiro, and Burke*, 1996], but in essence such models are only revisions to a curve fitting exercise. They may allow proton fluence to be predicted more accurately over timescales of years, but say nothing about *when* a specific SPE will occur.

2.4 The Flare Paradigm

The classical picture of SPE propagation to earth is based within a flare paradigm, and assumes particles to emanate from the location of a flare site where they then diffuse through space along magnetic field lines emanating from the sun. Burlaga describes a process for anisotropic diffusion in which particles are produced at a point source on the sun and diffuse to adjacent solar longitudes whilst at the same time diffusion proceeds more rapidly along interplanetary magnetic field lines away from the sun [*Burlaga*, 1967].

This diffusion process means that for a SPE to 'occur' at earth, solar protons must travel along magnetic field lines that connect the sun to the earth. More specifically, the diffusion model predicts that both the magnitude of solar proton flux seen at earth and the onset time are dependent on the heliolongitude of the particle source [*Barouch, Gros, and Massa*, 1971]. This has been verified observationally by a number of authors, most notably Shea and Smart, who show that the majority of SPEs seen at earth are related to flare sites on the western limb of the sun [*Shea and Smart*, 1995a], [*Smart and Shea*, 1989].

Cane et al. found that x-ray flares associated with SPEs could be grouped into two categories via their soft x-ray profiles (i.e. their profiles in the GOES 1-8Å channel), these being 'impulsive' x-ray flares, which displayed fast rise and decay times of the order of 10 minutes, and 'gradual flares' which exhibited a much longer exponential decay. Impulsive flares were shown to have small volumes and to occur low in the corona, with gradual flares occurring at much greater coronal heights. [*Cane, Mcguire, and Vonrosenvinge*, 1986]. Cane et al. found that particle emissions associated with impulsive flares tended to have a high ratio of electrons but generated a relatively low proton content and only occurred at well magnetically connected flare sites. By contrast SPEs associated with gradual x-ray flares had much higher proton energies, occurred at any solar longitude, and were also associated with coronal and interplanetary shocks.

The link between gradual x-ray flares and significant SPEs is corroborated in independent studies by Garcia, who identifies Gradual Hard X-ray flares as having a high correlation with SPEs [*Garcia*, 1994a], [*Garcia, Farnik, and Kiplinger*,]. Garcia also examined the electron temperature distribution of x-ray flares, and found that cooler flares with lower x-ray intensities had a higher association with SPEs, whereas hotter x-ray flares were not proton associated. Garcia theorised that gradual x-ray flares, as purported by Cane, belonged to this cooler, proton associated population, but also found that at high intensities the distributions merged: i.e. high temperature, high intensity flares could still result in associated particles [*Garcia*, 1994a]. At the same time, Kiplinger was examining the energy spectra of proton associated flares, and noted that x-ray flares that exhibited gradual hardening in energy spectra over their peak and decay were always associated with energetic SPEs [*Kiplinger*, 1995]. Kiplinger found that the behaviour of the energy spectra could be used to reliably predict whether or not an x-ray flare would result in associated protons, and theorised that flares with a hardening spectra were signatures of high energy electron and proton acceleration. Garcia and Kiplinger pooled their findings and found that large flares with SPE associations and progressively hardening spectra could have an impulsive stage of high electron temperature similar to that of non-proton associated x-ray flares which masked the characteristic low temperature usually seen in gradual proton associated flares [*Garcia and Kiplinger*, 1996]. Garcia and Kiplinger concluded that impulsive flare behaviour could occur just prior to a gradual flare event, with the effect of masking the typical gradual-flare temperature signature, and state that the fact that progressive hardening in energy spectra is still present suggests that a common acceleration process is present in all flares.

In summary Garcia and Kiplinger imply that SPEs are directly caused by energetic flares, and provide strong evidence that flare characteristics are good indicators of SPE occurrence. However, other findings indicate that coronal mass ejections (CMEs) may be the physical cause of SPEs rather than flares [*Reames*, 1995]. Kahler et al. found that 96% of energetic SPEs could be associated with the occurrence of a CME [*Kahler, Sheeley, Howard, Koomen, Michels, Mcguire, Vonrosenvinge, and Reames*, 1984], which are themselves highly correlated with the occurrence of gradual x-ray flares [*Kahler*, 1992]. This is also supported by Sheeley et al., who found that for a given x-ray event the probability of an associated CME is a strongly increasing function of x-ray flare duration [*Sheeley, Howard, Koomen, and Michels*, 1983]. The occurrence of a long duration x-ray flare is therefore a good indicator of both a CME and an energetic particle event, although there is no conclusive evidence as to which is the actual cause of a SPE. It is believed by some that flares play no physical role at all in causing SPEs, and that their supposed importance has been caused by the fact that they are relatively easy to monitor, (and have indeed been monitored since the 1930s). It is probably true that associations between SPEs and other phenomena have been overlooked due to the assumed role of flares, but Garcia and Kiplinger have proven that flare characteristics can be used to determine whether associated particles will be seen after a flare, and this will hold regardless of whether or not flares are the true cause of SPEs.

The relative particle abundances measured during SPEs indicate that there are probably two families of SPEs which emanate from two different regions of the corona, [*Reames*, 1988] and which relate to the two different types of flare association [*Reames, Barbier, Vonrosenvinge, Mason, Mazur, and Dwyer*, 1997]. Tanaka suggests that long duration x-ray flares are indicative of flares occurring high in the corona in areas of open and non-complex magnetic field topologies [*Tanaka*, 1987]. The mechanism of particle acceleration is thought to be due to magnetohydrodynamic shock wave propagation over large areas, and would explain the lack of correlation between gradual x-ray events and the sun-earth connection longitude. Impulsive flares with particle associations occur low in the corona and occur at well magnetically connected sites. Garcia postulates that if an impulsive flare is energetic enough it may allow particles to break their magnetic containment, giving rise to a low coronal source SPE. This concept is supported by Kahler's 'Big Flare Syndrome' which essentially states that any very energetic flare will have associated particles [*Kahler*, 1982b]. Although Garica and Kiplinger give evidence to suggest that the same (flare) acceleration mechanism is common in both impulsive and gradual cases, it is likely that the additional MHD wave acceleration process in the latter case is far more important in determining the severity of SPEs seen at earth.

It has also been accepted that interplanetary shocks play a fundamental role in SPEs by modifying the energy of particles that have been produced in earlier flares [*Cane, Reames, and Vonrosenvinge*, 1991], [*Cane, Reames, and Vonrosenvinge*, 1988]. This realisation stems from the observed behaviour of proton flux during large fluence events, which often exhibit several augmentations in proton flux associated with the passage of interplanetary shocks [*Kallenrode, Wibberenz, Kunow, Mullermellin, Stolpovskii, and Kontor*, 1993], [*Meyer, Wibberenz, and Kallenrode*, 1993]. Fundamentally this means that any model attempting to predict the flux of a proton event must also account for the effects of 'post-injection' interplanetary shock particle acceleration.

The argument as to whether it is CMEs or flares that cause SPEs is ongoing. Garcia and Kiplinger describe spectral hardening in x-ray flares as a manifestation of a particle acceleration process, providing strong evidence that flares are inextricably linked to SPEs, yet others infer that the high correlation between SPEs and CMEs must imply a physical link between SPEs and CMEs. However, in terms of wanting to predict the occurrence of a SPE it may not be necessary to debate their true physical cause. Garcia and Kiplinger have proven that flare characteristics can be used to determine whether associated particles will be seen after a flare, and this will hold regardless of whether or not flares are the true cause of SPEs. Purely from a forecasting viewpoint it is sufficient to identify and 'harness' precursors to SPEs so that forecasts can be made. Understanding the physical processes that cause SPEs is obviously desirable and may help to construct a forecasting model, but knowledge of these mechanisms is not mandatory to the success of a prediction technique.

2.5 Precursors and Associations

SPEs are nearly always co-incident with a solar x-ray flare, and as discussed above, it is widely held that particles are produced, (if not also accelerated) at the site of a flare. The radial emission of electromagnetic radiation at the speed of light as opposed to the diffusion of particles along open magnetic field lines gives rise to a delay between monitoring an x-ray flare and monitoring associated particles, hence flares are observed as precursors to SPEs. The delay between monitoring an x-ray flare and monitoring particles is variable, dependent upon the energy to which particles have been accelerated, and can vary between several hours, or just several minutes for very high energy particles (>500MeV) [*Smart and Shea*, 1989]. Flares however are a very common phenomena, with a typical occurrence frequency of several per day during solar maximum, the reality being that less than 1% of all flares produce associated particles at earth. The occurrence of a flare can therefore never be used as a unique indicator that an SPE will occur.

In addition to the flare characteristics described in the previous section radio bursts associated with flares have also been shown as having SPE associations [*Kahler*, 1982a]. Impulsive x-ray flares show strong association with Type III radio bursts, and gradual x-ray flares with Type II and Type IV radio drifts, which are indicative of a shock moving away from the sun , [*Cane, Mcguire, and Vonrosenvinge*, 1986], [*Aschwanden, Montello, Dennis, and Benz*, 1995]. Uddin et al. studied the radio associations from a sample of 52 proton associated flares and found that Type II, III, and IV radio emissions were present in 70% of cases [*Uddin, Pande, and Verma*, 1990]. Use of this correlation is made in current forecasting methods, but the fact remains that the presence or absence of a given radio burst type cannot be used as a unique requirement for a SPE. Uddin also finds for example that 38% of proton flares exhibit type I radio burst emissions.

It was mentioned in 2.3 that King modelled proton fluence as a function of sunspot number, but studies since 1970 have shown this to be a falsehood. Feynman et al. considers proton fluence between 1956 and 1985 and fails to find a relation between the solar cycle integrated proton flux and the maximum sunspot number of the cycle, and also on a year-by-year basis finds no correlation between the annual integrated solar flux and annual proton fluence [*Feynman, Armstrong, Daogibner, and Silverman*, 1990a]. Gabriel also shows poor correlation between sunspot number and annual proton fluence [*Gabriel, Feynman, and Spitale G.*, 1996] and a corroborative conclusion is also reached by Shea and Smart, who found that the annual occurrence rate of SPEs did not correlate well with the average annual sunspot number over solar cycles 19, 20 and 21 [*Shea and Smart*, 1992].

In an auto correlative approach Gabriel et al studied the occurrence frequency of SPEs by approximating SPEs occurring during cycles 19, 20 and 21 as unit pulses [*Gabriel, Evans, and Feynman*, 1990]. Whilst peaks in spectral power at near 154 and near 50 days were found, no level of significance was attributed to the findings, hence it is impossible to draw any conclusions. In addition, the tendency for events to occur in episodes of activity as an active region passes across the solar disk [*Shea and Smart*, 1990] means that SPEs from one initial source may have been included as separate events, where it may have been more appropriate to consider only the first injection. Nevertheless, the authors postulate that the observed 154 day periodicity may either be due to the correlation between SPEs and Hα flares, which have also been measured as having similar occurrence periodicity, or a more fundamental origin of solar activity connected to the periodic re-emergence of magnetic flux and modes of solar oscillation [*Oliver, Ballester, and Baudin*, 1998], [*Cane, Richardson, and Von Rosenvinge*, 1998]. As the significance of the periodicities cannot be quantified though it is impossible to take these theories as anything more than speculation.

Solar flares, and hence by association SPEs, are known to originate from active regions, consisting of local areas of complex, developing magnetic topology, which manifest as sunspots in visible light. Uddin et al calculated the cumulative flare index for 45 active regions that produced SPEs, and noted that a sharp rise in the cumulative flare index was often a precursor to SPE occurrence [*Uddin, Pande, and Verma*, 1990]. Importantly, the rise in cumulative flare index occurs over a timescale of days, implying that SPEs might be forecast without the need for a single x-ray flare, but Uddin makes no comparison to active regions that do not produce SPEs hence it is unknown as to whether sharp rises in cumulative flare index occur commonly when unrelated to SPEs. In a similar study Chakravorti does compare 171 proton to 200 non-proton producing active regions, and finds that the latter have a higher distribution of flare index values, but does not state how the flare indices have been collected: i.e. it is unknown if these are maximum values over active period life, or an average value [*Chakravorti, Das, Sen, and Dasgupta*, 1991]. Nevertheless, the result still implies that active region characteristics can be associated with SPEs. In addition, Chakravorti et al also find that the Ca-II plage index, the radio emission flux and the maximum intensity of 9.1cm radio flux associated with proton producing active regions pass through a maximum on the day of SPE occurrence in 70% of cases with values decreasing for days before and after. This observed behaviour again implies that SPEs might be forecast on a timescale of days.

In very recent work by Canfield et al. the appearance of Sigmoid shapes in x-ray images of the corona has been identified as a precursor to solar eruptions, and more specifically CMEs [*Canfield, Hudson, and Pevstov*, 2000]. MHD simulations show that the sigmoid forms when twisted magnetic fields combine in a reconnection mechanism, but the technique requires the magnetic structure to be seen 'end-on' for the sigmoid to be visible. Cleary this limits its usefulness as a practical precursor, and in addition it is not a unique pre-requisite for a CME, with only a 65% correlation with solar eruptions.

The current approach to SPE forecasting is to use x-ray flare characteristics to determine whether or not an x-ray flare is proton producing, but the findings by Chakravorti and Uddin described above strongly suggest that the characteristics of active regions (rather than discrete flares) may allow an SPE to be forecast a matter of days in advance, well before the traditional 'proton-flare' occurs.

2.6 Solar Proton Prediction

Probabilistic models such as JPL-91can provide an estimate with given confidence levels of the proton fluence that will be received over a number of years but do not address the immediate real-time hazards that SPEs pose because they do not predict when a specific event will occur. This means that SPEs currently have a large impact on real-time spacecraft operations and will pose a significant potential threat to future manned missions.

The current way to minimise the effects of SPEs is to design and use radiation hardened components, but this entails high costs, whereas an ability to alert operators to possible satellite errors could be more cost effective [*Shea and Smart*, 1998]. The threat posed by SPEs to future interplanetary missions is already appreciated [*Heckman*, 1993], and a reliable warning system would allow time for exploration teams to reach shelter or postpone a sortie. Even though protected by the earth's magnetic field, manned missions in Low Earth Orbit still experience higher radiation levels during SPEs, and again, the ability to forewarn of the occurrence of an SPE with suitable lead time and confidence could allow EVAs or even spacecraft launches to be rescheduled.

The immediate hazard that SPEs pose to spacecraft and manned missions requires the need for real-time forecasting models rather than probabilistic, long-term dose models. These real-time models need to be capable of predicting when an SPE will occur with enough lead time to take any mitigating action. The SEC has recognised the need for such real-time SPE forecasts and currently operates two models.

2.6.1 The PROTONS model

The PROTONS model entered operational service in 1972 and was developed in association with the Apollo lunar missions in order to provide an alert/warning service for energetic particles for manned missions [*Heckman*, 1988]. It was developed at a time when solar observations were still predominantly ground based and when solar mechanisms were poorly understood, and as a result the model is based on heuristic equations and empirical formulae derived from observed data ([*Heckman, Kunches, and Allen*, 1991], [*Balch and Kunches,*]. The PROTONS model is based firmly in the flare paradigm and is based on the concepts described previously in which energetic particles are produced at a flare site and then diffuse to and along interplanetary magnetic field lines connecting the sun and the earth.

The PROTONS model uses a solar flare to make two distinct predictions: the first is whether or not an SPE will occur, and the second is the >10MeV peak proton flux and arrival time of the particles should an event occur. When a flare occurs, PROTONS can be run, and takes as inputs the time integrated x-ray flux of the flare from the GOES instruments, the solar longitude and latitude of the flare, the current Ap index, and the presence of any type II and type IV solar radio bursts occurring in association with the flare.

The probability of a SPE following a flare is estimated by combining individual probabilities relating to different characteristics of the flare, thus the probability of there being a proton event takes the form:

$$\mathbb{P}(\mathrm{SPE}) = \mathbb{P}(1) \times \mathbb{P}(2) \times \mathbb{P}(3)$$

Where $\mathbb{P}(1)$ relates to the peak x-ray flux of the flare, $\mathbb{P}(2)$ relates to the type of radio burst associated with the flare, and $\mathbb{P}(3)$ relates to the magnitude of the predicted peak particle flux (a larger predicted proton flux indicates a higher probability that the proton event will actually occur). The probabilities for each term have been derived empirically and are based on experience with the model during 1974-1985. A proton event has a high probability of occurring if either Type II or Type IV radio bursts are observed, and the peak flux of the x-ray flare is high (greater than M6 level). Details of the probabilities used in the model are given by Balch [*Balch*, 1999]. Clearly, the probability of seeing an SPE after a flare has occurred is based purely on an empirical correlation between historical flare characteristics and SPE associations.

Predictions for expected peak proton flux are based on an empirically derived power law relating time integrated x-ray flux to observed peak proton flux. The value is corrected for the heliographic location of the x-ray flare, which attenuates the expected flux as a function of degrees from the sub earth point according to anisotropic diffusion theory.

The sub-earth point (helio-longitude) is itself a function of solar wind speed, for which the A_p index is used as a surrogate. Rise-time is also estimated from the flare location, and is a function of the distance between the flare site and sub-earth point.

Later additions were added to the model to account for the effects of interplanetary shocks giving rise to post injection particle acceleration. The modification however is relatively simple, and is again empirically based, simply stating that if a flare of sufficient size has occurred in the past 48 hours further correction factors are calculated based on the integrated flux of this flare. The effect is to increase the predicted particle flux and reduce the predicted rise time for the current SPE prediction.

PROTONS has been in real-time operation since 1972 which has enabled validation on a large amount of solar data, although the most recent verification and analysis was performed by Balch using 88 proton events with flare precursors occurring between 1988 and 1997 [*Balch*, 1999]. Balch finds that the predicted proton flux has an error of approximately one order of magnitude compared to the observed flux, and that the predicted lead times frequently vary by more than 50% from the observed lead times.

More fundamentally, Balch found that the flare location was not useful in predicting the peak flux of a proton event, and that the PROTONS model for peak proton flux prediction was marginally improved when the correction factor for flare heliolongitude was removed. Similarly, a slight improvement was also found when the correction for previous x-ray flare activity was removed. The significant errors in flux and lead-time predictions and the fact that ‘correction factors’ actually worsen performance reflect the simplistic diffusion model of particle propagation that PROTONS uses, and suggest that the empirical relations between x-ray flare energy and observed proton flux need to be revisited.

Balch also examined the accuracy of the model in predicting the occurrence of SPEs by comparing the predicted probability of there being an SPE with the actual observed frequency of SPEs, and found PROTONS to significantly over predict: e.g. model predictions of 40% and 50% for event probabilities actually had only a 21% chance of being associated with an event [*Balch*, 1999]. In an earlier study Heckman carried out an analysis of the model for 1989 in terms of a 'yes/no' forecast derived by taking all probabilities > 50% as a 'yes' and all those < 50% as a 'no'. During 1989 only 5 out of 21 events were missed, resulting in a hit-rate of 76%, but in 16 cases an SPE was predicted when none occurred, causing Heckman to state that the model generates a relatively large proportion of false alarms [*Heckman, Kunches, and Allen*, 1991].

PROTONS is a model which relies on relatively simple physics and empirical relations between x-ray flares and SPEs. In terms of predicting whether an SPE will occur or not, performance is reasonable, but the quantitative outputs such as occurrence probability, peak proton flux and rise time have significant error, reflecting the simplistic physical origins of the model.

2.6.2 The Garcia Model

The Garcia model again exists under the flare paradigm but is based on the more recent findings discussed in 2.4 in which Garcia found proton producing x-ray flares to have a low temperature distribution at low peak fluxes [*Garcia*, 1994a]. Garcia fitted probability curves to the observed distributions, allowing the probability of an SPE occurring to be calculated from flare electron temperature and peak x-ray flux.

This model is currently being evaluated on real-time data on SEC web space♠ using x-ray fluxes and flare temperatures calculated from GOES satellite data [*Garcia*, 1994b] and has yet to be formally evaluated. However, Garcia shows that the two distributions of low and high temperature merge at high fluxes, indicating that in some instances the model will not be able to distinguish between proton and non-proton flares.

It is interesting to note that Kiplinger's findings, which were also discussed in 2.4 relating to the evolution of spectral hardness during x-ray flares, have not been used to create a particle prediction model, even though Kiplinger performed a test in which he was able to successfully identify 96% of proton producing flares purely through analysis of their x-ray spectra [*Kiplinger*, 1995]. The most probable reason for this is that the spectral hardening of the x-rays must be examined over the decay period of the flare, by which time protons associated with the flare may already be arriving at earth: i.e. the technique only permits a post-event association to be made and cannot be used as a practical prediction method.

♠ http://sec.noaa.gov/~sgreer/gprot/index.html

2.6.3 The Problem with Flare Paradigm Prediction Models

It is evident that the only particle prediction models that are in current operation rely on the characteristics of a single, discrete x-ray flare in order to make a prediction. The approach taken is to look at an x-ray flare and decide whether or not it is going to be proton producing. Whilst there may be a physical argument against using flares as a predictor because they are not the true physical cause of particles (as discussed in 2.4), there is also a more fundamental, practical drawback to the technique in that particles can arrive just several minutes after the occurrence of a proton flare [*Smart and Shea*, 1989]. Typically, particles will arrive within several hours of an associated x-ray flare, meaning that the PROTONS, Garcia and any other technique using flare characteristics will always have lead times that are physically limited to just a few hours, and in some instances, for very energetic particles (which are arguably the most desirable to be able to forecast), there will be almost no lead time at all.

Secondly, flare paradigm models do not allow an SPE to occur without there first being a flare, whereas real observations show that it is possible for an SPE to occur without an apparent x-ray flare. In Heckman's appreciation of the performance of PROTONS during 1989 he finds that a proton event was not predicted because it occurred without a flare [*Heckman, Kunches, and Allen*, 1991], hence there was no reason (and no inputs) to run the model. Whilst these instances are relatively uncommon it still highlights another inherent limitation with current particle forecasting methods.

Future interplanetary missions or human expeditions to Mars or the Moon are likely to require some kind of solar particle warning service, but the current lead times, and accuracies that are offered by such models may not be sufficient to allow an acceptable level of risk [*Heckman*, 1993]. A recent work commissioned by the European Space Agency set out to quantify the current Space Weather requirements within the European Community and identified a need to predict solar particles for satellite operators and manned missions with forecasts of 1 to several days ahead with target reliabilities of 65% [*Horne R. B.*, 2001].

The SEC does issue the probability of an SPE occurring for 1, 2 and 3 days ahead, but these probabilities are generated by a human forecaster and not by an analytical model. Forecasts are based on a number of factors including magnetic complexity of active regions, active region area, H-alpha and white light structure, heliographic position of active regions, coronal structure, radio emission characteristics and whether or not active regions are in growth or decay [*Heckman*, 1993]. These predictions are intended to serve as a general outlook, but analysis by Heckman shows that the predicted probabilities bear almost no relation to the observed occurrence of SPEs: i.e. these forecasts have no reliability whatsoever.

2.7 A New Particle Prediction Model

It has been shown that current SPE forecasting models are deficient. PROTONS exhibits significant error in rise time and peak flux predictions, and the underlying physics on which it is based is overly simplistic. Predictions of whether or not an SPE will occur are of reasonable reliability, but the pre-requisite of an x-ray flare by both PROTONS and the Garcia model limits lead-times and treats SPEs uniquely within a flare paradigm. The need to forecast SPEs several days ahead has been identified, yet the flare-based models cannot achieve this, and alternative forecasting techniques are subjective and highly inaccurate on such a time-scale.

There is evidence however to suggest that precursors to SPEs other than the classical solar flare may be present, and may allow events to be forecast on a timescale of days. As was discussed earlier in 2.5, Uddin et al found solar radio flux from active regions to pass through a maximum on the day of an SPE, yet these findings have not been taken further: the behaviour of other solar variables, such as x-ray flux, has not been studied on a timescale of days relative to SPEs, (possibly due to the well established belief that flares/CMEs are the first and only precursors to SPEs) and importantly, no attempt has been made to create an SPE forecasting model that does not use a discrete x-ray flare as an input.

The existence of several decades of in-situ high quality solar measurements provides an excellent opportunity to address the 'gap' in the current SPE prediction approach by providing an opportunity to try and identify longer term precursors to SPEs. Via this approach it may be possible to develop a prediction model that does not require characteristics of a discrete flare as inputs, hence permitting the generation of forecasts with longer lead times. In addition the wide availability of near-real time data from satellites provides an excellent resource from which to evaluate and operate such a model. If successful, such a technique would go some way to realising the particle forecast lead-time requirements that have been identified by ESA and which would be required by future manned space missions.

3. A NEW PREDICTION MODEL

This chapter examines the nature of SPE occurrence with reference to specific examples. By examining the limitations of current prediction techniques it sets requirements for a new forecasting model and different approaches are then discussed as to how these requirements can be met.

3.1 The Nature of SPEs

In order to begin the development process for a new SPE prediction technique it is useful to examine the occurrence of an SPE in order to characterise the phenomenon in more detail. Since the existence of space borne particle monitors SPEs have been observed via the instruments of the IMP and GOES satellite platforms. The Space Environment Monitor payload on board the GOES satellites also contains solar x-ray detectors, allowing solar x-ray activity to be observed in parallel with particle fluxes. Measurements from GOES are available in near real-time as a space weather monitoring service, and are the means today by which SPEs are most commonly observed and analysed.

Figure 3-1 shows an SPE which occurred at near mid-day on the 14th February 1986 as recorded by monitors on the GOES-6 spacecraft. Prior to the SPE the >10MeV proton flux is at a relatively constant level of approximately 10^{-1}p.f.u. The SPE manifests as a sudden increase in flux over two hours to a level of ~10^{1}p.f u, followed by a more gradual increase to approximately 10^{2} p.f.u. over a further 6 hours. The elevated flux levels then decay to background over a period of 3 days (although the tail of the event is not shown in the figure). An examination of the solar x-ray flux shows an x-ray flare to occur just prior to the onset of the SPE. The time delay between the peak of the x-ray flare and the onset of the SPE is approximately 2 hours.

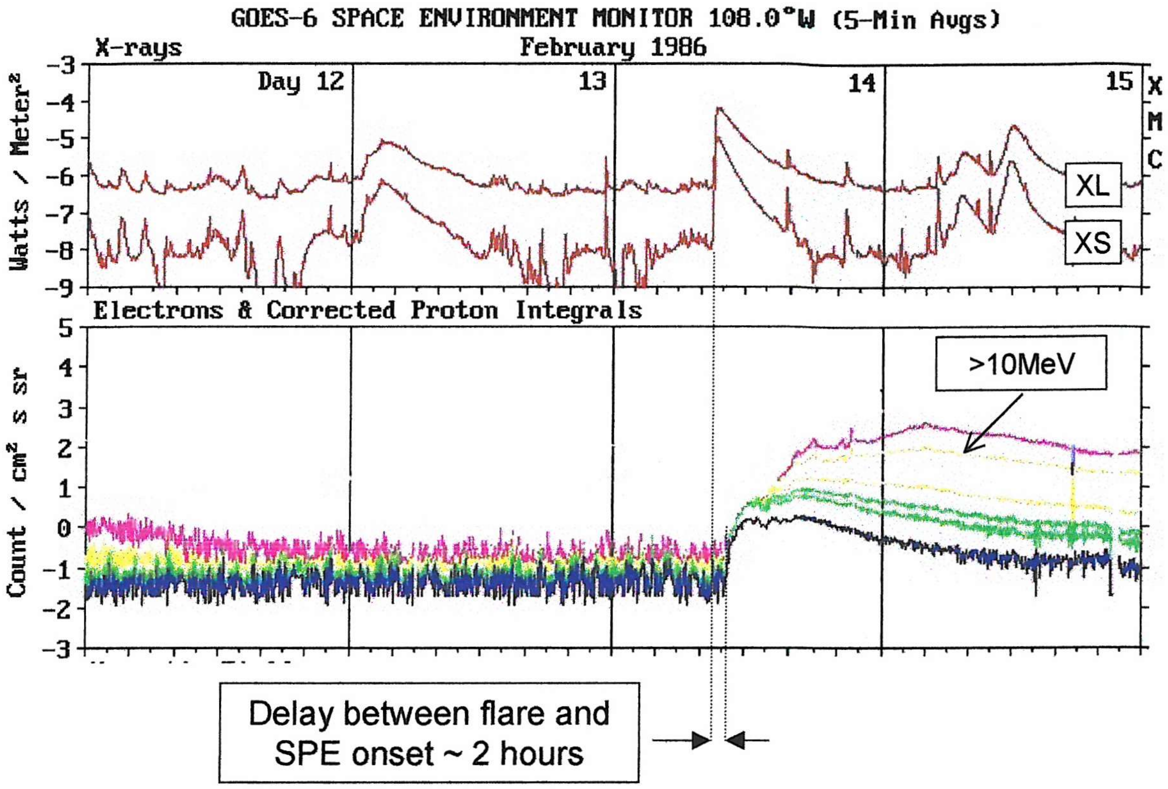

Figure 3-1 Typical occurrence of an SPE with associated x-ray flare during February 1986 as captured by monitors on the GOES-6 spacecraft.

Figure 3-2 and Figure 3-3 show similar plots for two more arbitrarily chosen SPEs, one from November 1988, the other from July 1989. In both instances the >10MeV proton flux is again at a relatively constant background level of ~10^{-1}p.f.u. prior to the event. The SPE manifests as a sudden sharp increase in the proton flux level. In the July 1989 example (Figure 3-2) the SPE is associated with an impulsive x-ray flare of > 10^{-4} watts/m^2 peak flux, but there is no measurable delay between the x-ray flare peak and the onset of the SPE. In the November 1988 example the onset of the SPE does not appear to have any x-ray flare association at all.

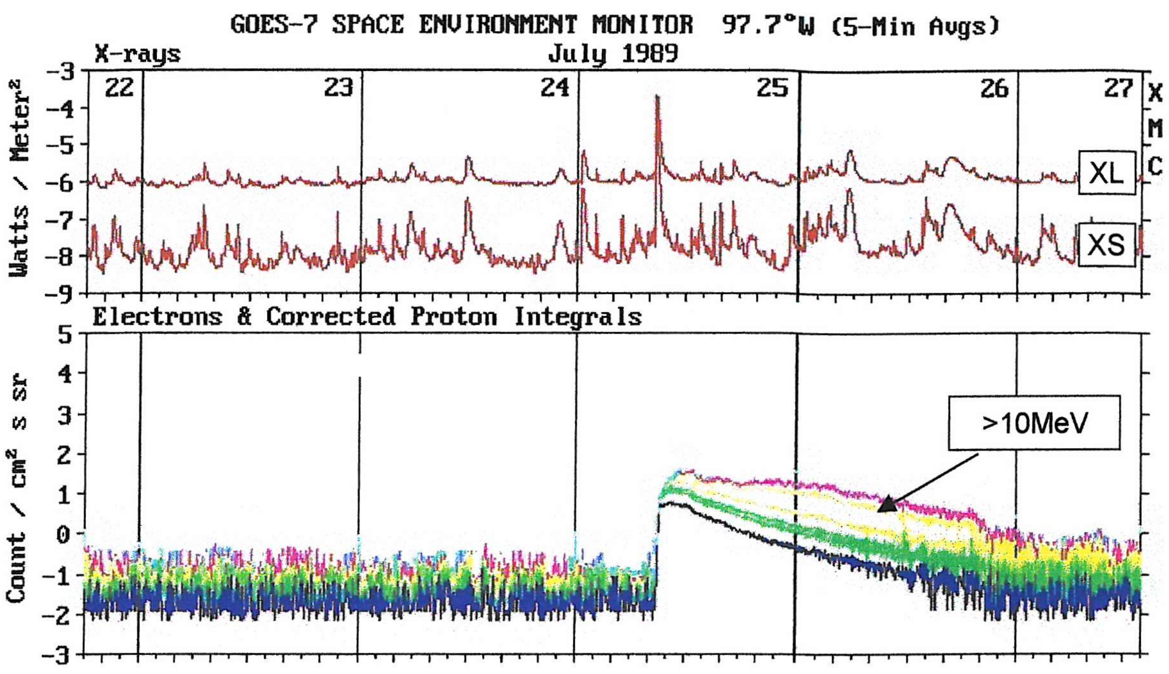

Figure 3-2 An SPE with associated x-ray flare on 25th July 1989

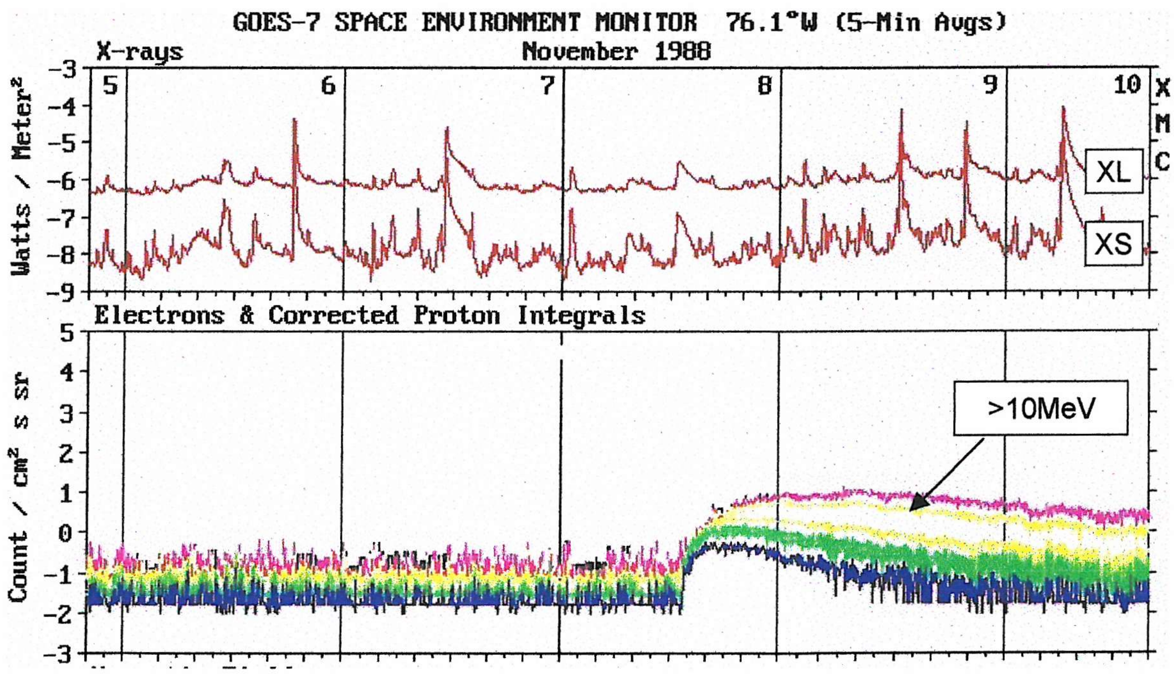

Figure 3-3 an SPE with no associated x-ray flare on November 8th 1988

The specific flux-time profile of SPEs, as discussed in the previous chapter, is a function of the heliographic position of the particle source with respect to the earth and also dependent on the presence of interplanetary shocks, but in the general case all SPEs are similar in that they are typified by a sudden enhancement in proton flux which rises to a maximum over a period of several hours. Prior to an SPE the >10MeV proton flux is at a relatively constant background level of ~ 10^{-1}p.f.u.

3.2 The Current Forecasting Approach

The current approach to SPE forecasting is to determine whether or not an x-ray flare is likely to produce protons by examining characteristics of the x-ray flare. Analysis of the PROTONS model during 1989 by Heckman shows that the technique makes a correct forecast in 78% of cases [*Heckman, Kunches, and Allen*, 1991], hence the accuracy is actually very reasonable, but the technique has two fundamental drawbacks. These are the lead-time of predictions and the requirement that an x-ray flare must occur for an SPE to occur.

For the PROTONS model, the lead time for an SPE prediction is a function of flare location and energy. Based on this information the model predicts a rise time for the SPE, which is the time between the x-ray flare peak and the maximum >10MeV proton flux of the SPE. Typically an SPE will take between 1 and 24 hours to reach its maximum flux, but in practice, as can be seen from the previous examples, significant particle fluxes are still observed within 0-6 hours of flare occurrence, well before the peak proton flux is reached. This means that the PROTONS model can effectively only produce forecasts a few hours before an SPE begins. The Garcia model suffers from exactly the same problem in that an x-ray flare has to peak before a prediction can be made (the model uses x-ray flare temperature and peak flux in order to make a forecast).

The fact that existing models use flare characteristics in order to make predictions also means that an x-ray flare must occur for an SPE to be predictable. Based on the dataset used later in this study, 93% of SPEs occur in association with an x-ray flare, so in practice flare based are justified, but fundamentally it still means that an SPE can occur without there being any inputs on which to run a forecasting model. Clearly, flare based SPE prediction models are incapable of predicting SPEs that occur without an x-ray flare.

3.3 The Requirements of a New SPE Forecast Model

It has been shown that flare based SPE forecasting models can give predictions with reasonable accuracy, but as they need to wait for a flare to occur and peak before they can make a forecast, significant proton fluxes can be seen within minutes of a forecast being made. A new prediction approach should address this by producing forecasts with lead times of the order of a day or more (i.e. 24-hours before significant proton fluxes are seen). Such a period would allow a realistic time for any mitigating action to be taken, especially in the case of manned missions. For example, intensive spacecraft operations could be postponed rather than risk incurring an anomaly part way through a procedure that might leave a spacecraft in an undesirable configuration. In addition, sensitive equipment could be turned off completely rather than risk spurious commanding or degraded operation. In the case of manned surface exploration of the moon astronauts would require time to reach a radiation shelter.

Integral to the operation of an SPE forecasting model is the ability to operate in real time meaning that any inputs that are used must be available with a minimum lag from real-time. Inputs that require a significant amount of time to calculate or retrieve will eat into the lead time of the model and reduce its capability. The requirements of a new SPE forecasting model are summarised in Table 3-1. The exact outputs from a new prediction model have not been stipulated, and will in part depend on the modelling technique that is employed.

Requirements for a new SPE prediction model
• Provide lead times greater than current flare association models (24 hours +)
• The model must be capable of real time operation (i.e. input variables must be available in real-time or near real-time)

Table 3-1 : Requirements for a new SPE prediction model.

The aim to produce forecasts with >24hours lead time implies that discrete x-ray flares cannot be used as inputs to a model, hence any new model will be a departure from the classical flare association models that currently exist. It is important to note that to date, no attempt has been made to predict an SPE without requiring an x-ray flare, hence the development of a 'long' lead time model will be the first investigation of its kind.

It is also important to define what is meant by SPE prediction. For the basis of this study the goal is to predict the sudden onset of an SPE. The prediction of quantitative characteristics of an SPE, such as peak flux and total fluence are not objectives.

3.4 Alternative Forecasting Approaches

The requirements for a new prediction model can potentially be met by a number of methods. One broad method is via a physics based model whereas another is via the empirical route relying on correlations between SPEs and other observed variables. Potential techniques are discussed below.

3.4.1 Physical Model

A physical model attempts to describe a process by understanding the basic mechanisms that govern it. Such an approach requires the physical processes to be identified and understood to such a level that they can be modelled mathematically from first principles and driving variables need to be identified so that they can be used as inputs to the model.

In terms of modelling solar proton events there are several potential different mechanisms that would need to be considered, corresponding to different stages in SPE occurrence. Speculatively these might be classified as a pre-injection phase, in which active regions are evolving, an acceleration stage, in which particles are injected and accelerated within the corona, a solar propagation stage, in which particles travel through the solar atmosphere, and an interplanetary propagation stage, in which particles travel along interplanetary magnetic field lines and are potentially re-accelerated by further shocks. Table 3-2 summarises these potential mechanisms, tentatively identifies their physical processes, and speculates as to what the required input variables might be for the relevant models of each process.

Mechanism to be Modelled	*Processes involved*	*Possible Input Variables*	*Related References*
Development of active region complexity	Magnetic Evolution	Magnetic field strength Solar Rotation Helio lat/long	[*Reames*, 2002;*Cline, Brummell, and Cattaneo*, 2003;*Kalman*, 2001;*Akasofu*, 2001;*Priest and Forbes*, 2002]
Shedding of magnetic structure	Thermal Magnetic	Magnetic field strength Solar Rotation Helio lat/long	[*Bieber and Rust*, 1995;*Moore*, 1988]
Impulsive particle acceleration	Thermal	Driving energy Particle population	[*Wu, Wang, Xu, and Tang*, 2002;*Miller*, 1998]
MHD Shock particle acceleration	Thermal Magnetic	Driving Energy Particle population CME characteristics	[*Vrsnak and Lulic*, 2000a;*Vrsnak and Lulic*, 2000b]
Particle propagation around the sun/corona	Particle Diffusion	Helio lat/long Magnetic Field Topology Solar Wind Speed	[*Balch*, 1999;*Kunches and Zwickl*, 1999;*Huang and Wang*, 2001]
Particle propagation through interplanetary space	Particle Diffusion	IP Magnetic Field Solar Wind Speed Post acceleration particle energies	[*Cargill, Chen, Spicer, and Zalesak*, 1996;*Chen and Garren*, 1993]
Interplanetary shock acceleration of particles	Thermal	IP Magnetic Field Particle Energies IP CME Propagation	[*Reames*, 1990;*Kallenrode*, 1995]

Table 3-2 Possible mechanisms and inputs to be incorporated in a physical SPE prediction model.

An intrinsic problem with the physical approach is that the arrival of an SPE at earth is the culmination of a number of linked physical processes which are all complex in nature. Input variables at each stage of the model may need to be measured directly, or they may be outputs from a preceding stage of the model, in which case errors may propagate throughout the entire modelled process.

It is important to note that the PROTONS model already uses anisotropic diffusion theory to model the transport of particles around and away from the sun, yet as discussed previously, results are in significant error to observations because the presence of interplanetary shocks is not physically modelled at all [*Balch*, 1999]. This indicates that the physical processes need to be modelled accurately rather than approximated by simpler theory. Similarly, there is evidence to suggest that the presence of coronal holes situated between particle injection sites and the sun-earth connection point modulates the particle flux seen at earth [*Kunches and Zwickl*, 1999], suggesting that it may be necessary to model the whole of the solar magnetic field before accurate predictions about particle fluxes at earth can be made. The physical approach is confounded further by the fact that some of the physical mechanisms for particle acceleration are still being argued, and hence cannot be modelled with any certainty [*Reames*, 1999], [*Reames*, 1990].

Even if a physical model were to be created that did accurately describe particle acceleration and propagation, it still may not be capable of SPE prediction. SPE prediction requires the acceleration process itself to be predicted, which in turn means that the dynamics and energy of the suns surface magnetic field must be measured and modelled to such a degree that accurate extrapolations concerning its future evolution can be made. Essentially this means that flares themselves must be predicted (presumably from magnetic field configurations). Whilst such a task may not be impossible, it is likely to require high resolution measurements of active regions and the use of MHD simulations to model multiple, complex magnetic structures. It is highly ambitious to expect to achieve an integrated prediction model with real-time capability using such an approach.

3.4.2 Time Series Prediction

A time series prediction technique attempts to predict future values of a time series from past values of the same time series. One class of time series model is the Autoregressive Integrated Moving Average (ARIMA) model, which has been developed most notably by Box and Jenkins [1976]. The technique is powerful, being applicable to non-stationary series, and is capable of producing accurate forecasts. It also benefits from having a structured, formal approach that is well understood and documented.

The ARIMA approach as applied to SPE forecasting will treat the problem as one of time series prediction in which past values of proton flux are used to forecast future values. This means that output from the model will consist of quantitative values of proton flux. This is an inherent benefit of the technique, but at the same time, such a model can only be successful if the sudden enhancements in proton flux due to SPEs are actually predictable from the proton flux itself. There is no documented evidence of SPE precursors being present in the proton flux, and the earlier examples in Figure 3-1 to Figure 3-3 would appear to support this, showing the proton flux prior to SPE occurrence as being at a relatively unchanging background level.

However, despite the fact that the proton flux is not thought to contain SPE precursors, ARIMA models do provide a powerful approach to forecasting problems in general and, providing that a series is deterministic to some degree, allow accurate forecasts to be made. Although one might not expect SPEs to be predictable from a time series approach, the benefits of a formal ARIMA technique mean that if such a model were successful it would be relatively easy to implement operationally. In addition, the large amount of historical proton data that is readily available, from GOES satellites for example, also means that obtaining sufficient measurements on which to base an ARIMA model is not an issue, hence a time-series model can be created without need for a large investment in resources.

3.4.3 Alternative Classification Approach

An alternative approach to SPE prediction is to try and relate the behaviour of solar variables to the occurrence of a proton event. In principle the concept is not dissimilar to that of existing flare association models, but the important difference is that inputs would not be restricted to discrete x-ray flares. Inputs would be taken over longer periods of time, days before an SPE occurs.

Such a technique could potentially harness the longer term SPE related behaviour reported by Chakravorti et al, who found the solar radio flux from active regions to rise prior to an SPE and pass through a maximum on the day of an SPE. The behaviour of other solar variables, such as the solar x-ray flux, could be considered over a similar time period, (although the inputs to any such operational model would be restricted to data that is available in real-time, or near real-time).

The development of a correlative approach would require a list of SPEs in order that data pertaining to their occurrence times can be obtained. This data could then be perused in order to try and quantify any precursive behaviour to an SPE, with the aim of establishing the basis for a forecasting technique.

3.5 Selection of a Forecasting Approach

The previous section has considered three broad approaches that could be adopted in order to generate a new SPE forecasting model that meets the defined requirements.

The physical approach, whilst perhaps having the highest potential for accurate forecasting, probably has the lowest chance of success. The effort required to understand and simulate some of the key processes in SPE production means that emphasis will be placed on modelling a process and not on producing a real-time prediction tool. Given the current understanding of the underlying physics that govern SPE's and solar flares it is highly unlikely that a physical model can be designed, built and validated within the envisioned time frame.

The time-series prediction approach offers potential gains in terms of accuracy and the benefit of a formal method that is well understood. The likelihood of the technique succeeding is probably quite low as there is no evidence to suggest that the proton flux contains precursors to SPEs, but as the ARIMA technique has never been applied to a solar proton flux time series before the result will have value regardless of its success and provide a good opportunity to characterise the proton flux time series further. An ARIMA technique represents a good point at which to begin forecast model development.

A classification approach would attempt to correlate the behaviour of solar quantities with the occurrence of an SPE and the development of the required dataset also provides an opportunity to look for the existence of SPE precursors in different solar variables. However, implementation of an ARIMA technique as a first step will allow the proton time series to be analysed in more detail and may help implementation of a classification approach at a later stage.

4. TIME SERIES PREDICTION OF SOLAR PROTON FLUX

This chapter performs a time series analysis of the >10MeV solar proton flux over solar cycle 22. Background to the ARIMA modelling technique is given, followed by its application to the observed proton flux time series. ARIMA models are then assessed and conclusions drawn as to their accuracy and suitability to the problem.

4.1 ARIMA Time Series Modelling

The ARIMA technique assumes a time series to be composed of a deterministic component and a non-deterministic component (i.e. a random noise contribution). Mathematically, the value of a time series at time t (x_t) can thus be written as:

$$x_t = \mu_t + \varepsilon_t \qquad \textbf{4-1}$$

where μ_t is the deterministic mean of the process at time t and ε_t is a random error component.

ARIMA models are based on the supposition that successive observations within a time series can be represented by a linear combination of independent variables derived from previous values in the time series. In the simplest case a white noise process can be represented by independent random variables that are drawn from a normal probability distribution of mean 0 and variance σ, which can be written as:

$$x_t = \mu + \psi_0\varepsilon_t + \psi_1\varepsilon_{t-1} + \psi_2\varepsilon_{t-2} + \ldots \qquad \textbf{4-2}$$

where x_t is the value of the series at time t, ψ_j are constants, ε_{t-j} are the independent variables and μ is a constant determining the average of the process. This form of time series model is a linear filter, and is the base form of model which Box and Jenkins have developed into the ARIMA class of models [1976]. A time series can be adequately modelled using relatively few parameters of the infinite series.

A special case of the linear filter, is the autoregressive process in which the current observation, x_t, is regressed on previous values of the series: $x_{t-1}, x_{t-2}, \ldots, x_{t-p}$. This can be expressed mathematically as:

$$x_t = \xi + \phi_0 x_t + \phi_1 x_{t-1} + \phi_2 x_{t-2} + \ldots + \phi_p x_{t-p} + \varepsilon_t \qquad \text{4-3}$$

and is termed an autoregressive process of order 'p', abbreviated AR(p), where ξ is a constant and ε_t a random noise component.

Consideration of the linear filter (4-2) with only the first q weights non zero results in the other finite process:

$$x_t = \mu + \varepsilon_t - \theta_0 \varepsilon_t - \theta_1 \varepsilon_{t-1} - \theta_2 \varepsilon_{t-2} - \ldots - \phi_q \varepsilon_{t-q} \qquad \text{4-4}$$

where θ_j are weights, μ is a constant and ε_{t-j} are independent variables representing the random component at time t-j. This is a moving average process of order q, abbreviated MA(q).

In constructing the model for a time series the inclusion of both autoregressive and moving average terms can lead to a more efficient and accurate model than using either of the two forms alone. This results in the combined autoregressive-moving average model (ARMA) model of order (p,q):

$$\begin{aligned} x_t = \xi + \phi_1 x_{t-1} + \phi_2 x_{t-2} + \ldots + \phi_p x_{t-p} \\ - \theta_1 \varepsilon_{t-1} - \theta_2 \varepsilon_{t-2} - \ldots - \phi_q \varepsilon_{t-q} + \varepsilon_t \end{aligned} \qquad \text{4-5}$$

The application of the AR, MA and ARMA processes is applicable only to stationary time series, stationarity meaning that the joint distributions of any set of X_{t1}, X_{t2}, ….., X_{tk}, … are unchanged if the times t1, t2, ….tk are shifted by some value 's'.

In practice, time series may be non-stationary in mean, or non stationary in mean and slope, but can be made stationary by deriving a time series from the difference between consecutive terms in the series. A non-stationary time series can sometimes be reduced to a stationary series by applying a suitable degree of differencing. Differencing can be expressed using the difference operator ∇ and is shown in equations 4.6 and 4.7 for the 1st and 2nd differences:

$$\text{1}^{st}\textit{ difference:}\quad \nabla x_t = x_t - x_{t-1} \qquad \text{4-6}$$

$$\text{2}^{nd}\textit{ difference:}\quad \nabla^2 x_t = x_t - 2x_{t-1} - x_{t-2} \qquad \text{4-7}$$

Time series models of non-stationary series can sometimes be constructed by substituting x_t with a suitable difference term ∇^d in equation 4.5 to produce an autoregressive integrated moving average process (ARIMA) of order (p,d,q). Such a model represents the d^{th} difference of the original non-stationary time series as a process containing p autoregressive and q moving average parameters.

4.1.1 Determining an ARIMA model

A first approximation to the form of ARIMA time series model for a given series is made through analysis of historical data, in particular by examining the autocorrelation function of the sample, defined as:

$$r_k = \frac{\sum_{t=1}^{N-k} (x_t - \bar{x})(x_{t+k} - \bar{x})}{\sum_{t=1}^{N} (x_t - \bar{x})^2} \qquad \text{4-8}$$

Where r_k is the k^{th} autocorrelation co-efficient, and N the length of the sample time series. The autocorrelation is the correlation between the original series, and a copy of the series displaced by a lag of 'k' time intervals The partial autocorrelation function is also useful in determining the form of the model, and can be estimated from the autocorrelation coefficients. It is defined in 4-9.

$$\phi_{kk} = \frac{r_k - \sum_{j=1}^{k-1} \phi_{k-1_j} r_{k-1}}{1 - \sum_{j=1}^{k-1} \phi_{k-1_j} r_j} \qquad 4\text{-}9$$

The sample partial autocorrelation function is the correlation between the elements x_t and x_{t-k} on the set of intervening values $x_1, x_2 \ldots x_{t-k+1}$. It measures the dependence between x_t and x_{t-k} after the effect of intervening values has been removed.

Plotting the autocorrelation and partial autocorrelation functions allows the behaviour of the sample's functions to be compared with theoretical autocorrelation patterns of AR and MA processes (shown in Table 4-1). AR and MA components of a time series model can then be chosen according to those that are best suggested by the observed autocorrelations.

Model	Autocorrelation Function	Partial Autocorrelation Function
AR(p)	Tails off	Cuts of after lag p
MA(q)	Cuts of after lag q	Tails off
ARMA(p,q)	Tails off	Tails off

Table 4-1 Behaviour of theoretical Autocorrelation and Partial Autocorrelation Functions for AR and MA processes.

After identification of the form of ARIMA process, least squares estimates of the model coefficients are found by regression with the sample data. The adequacy of the model can be assessed by examining the residuals which should resemble white noise (i.e. a random error with zero mean) if the model is correct. The autocorrelations of the residuals should be zero for all lags greater than one. Non zero correlation coefficients indicate that the model may benefit from additional AR or MA terms [*Montgomery D.C., Johnson, and Gardiner J.S.*, 1990].

4.2 Method

An ARIMA model has been constructed based on the daily averaged >10MeV integrated proton flux measured by GOES satellites over the active period of solar cycle 22 (corresponding to 26/5/1987 to 25/5/1994). A plot of the time series is given in Figure 4-1. A daily averaged time resolution has been used in order to generate a time series that incorporates several SPEs whilst keeping a manageable size of dataset. It can be seen from the Figure that the proton time series is bursty in nature, with SPEs manifesting as sporadic enhancements in the flux level.

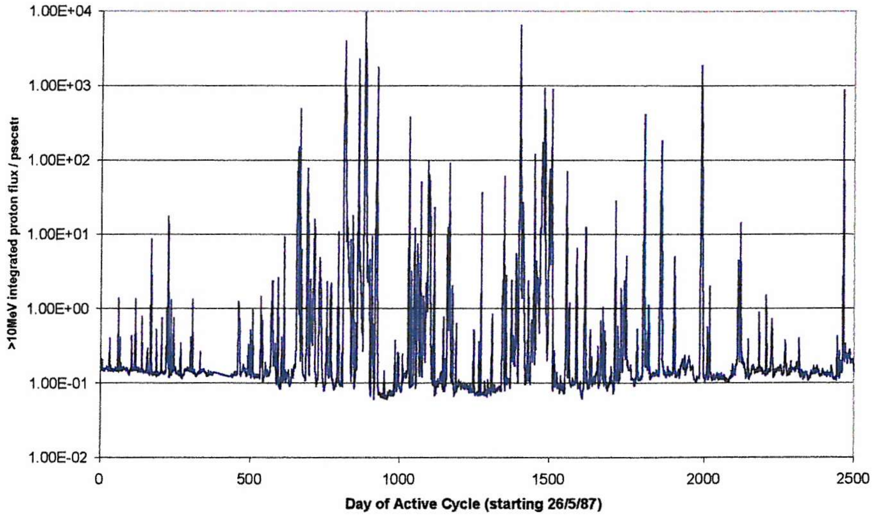

Figure 4-1: Variation of >10MeV Integrated proton flux during active period of solar cycle 22.

Data was checked for erroneous or missing values but none were found in the daily averaged GOES data for the period in question. This reflects the high resolution and quality of data available from the GOES satellites.

A sample of the time series was taken on which to base the ARIMA model, consisting of the first 1200 days of the series. Literature states that ARIMA models should be based on at least 60 measurements, hence this criteria is more than satisfied by the sample length [*Montgomery D.C., Johnson, and Gardiner J.S.*, 1990]. It can be seen from Figure 4-1 that the first 1200 days of the series contain several enhancements of greater than 1.0 p.f.u., which would be recognised as SPEs under the SEC definition.

In the first stage of analysis the sample time series was assessed via examination of the autocorrelation function in order to determine stationarity. A suitable degree of differencing was then applied to achieve stationarity.

The autocorrelation functions for the differenced time series were then examined in order to determine a form of ARIMA model. Table 4-1 summarises how the behaviour of the autocorrelation function corresponds to theoretical AR and MA processes.

Coefficients for the ARIMA models were calculated using an iterative routine in MINITAB (a commercially available statistics software package), and different forms of the model were generated in order to provide a comparison. The accuracy of models was assessed using the normalised root mean square error, defined as:

$$NRMSE = \sqrt{\frac{\sum_{i}^{n}(x_i - y_i)^2}{n\sigma^2}}$$

where 'n' is the total number of examples in the test set, x_i is the actual ith value, y_i is the predicted ith value and σ is the standard deviation of the actual values from their mean. A NRMSE of 1.0 is equivalent to using the mean of the data as a predictor whereas a NRMSE of 0 indicates a perfect fit.

In its standard form an ARIMA model will predict one time-step ahead, but forecasts of more than one time step can be made by running the model with predicted values as inputs. Forecasts were thus generated for 1, 2, 3 and 4 days ahead and compared against NRMSE values for persistence.

4.3 Results

4.3.1 Stationarity transformation

The stationarity of the sample time series was assessed by examining the autocorrelation function, which should rapidly die away to zero if the series is stationary. The standard error of the autocorrelation function was used to identify significant non-zero terms.

Figure 4-2 plots the autocorrelation function, calculated using MINITAB, for the original sample time series. The autocorrelation function does not die away rapidly, indicating that the time series is non-stationary. Taking the 1st difference of the series was found to produce stationarity, as can be seen from the behaviour of the respective autocorrelation function in Figure 4-2.

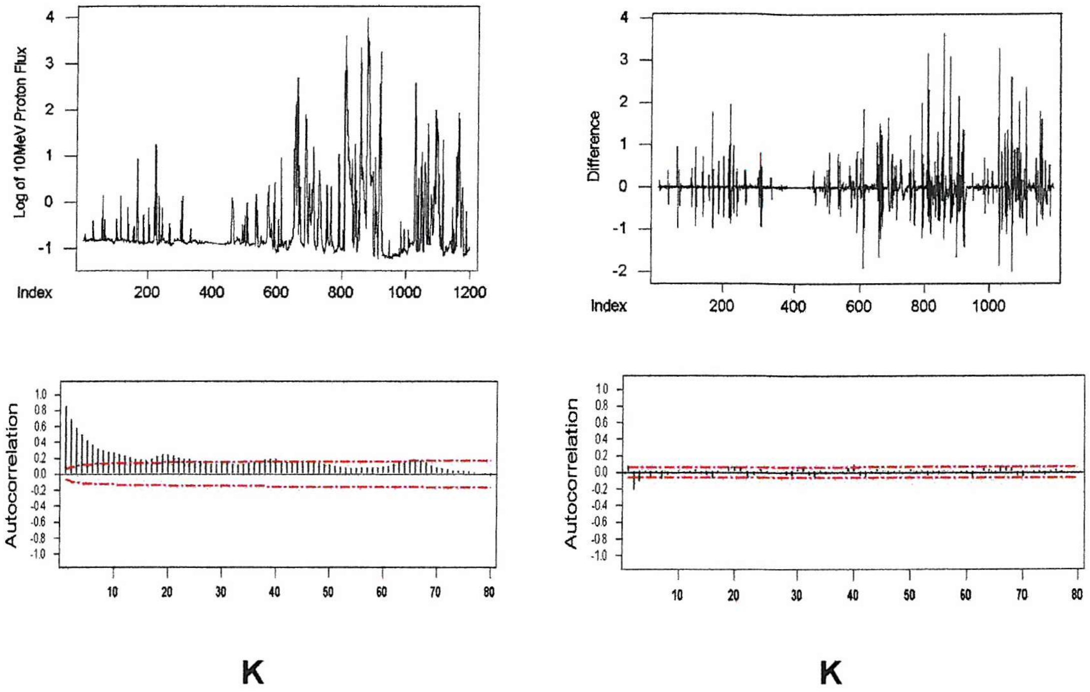

Figure 4-2 Time series plot and Autocorrelation function for >10MeV proton flux time series and the 1st difference of the original time series. The standard error of the autocorrelation functions is plotted as a limit in order to determine significantly non-zero terms.

4.3.2 Derivation of ARIMA model

Examination of the autocorrelation and partial autocorrelation functions with respect to the behaviour detailed in Table 4-1 was used to infer a form for the ARIMA model. The autocorrelation functions for the 1st difference are plotted below in Figure 4-3. Standard error margins have been used in order to identify significantly non-zero terms.

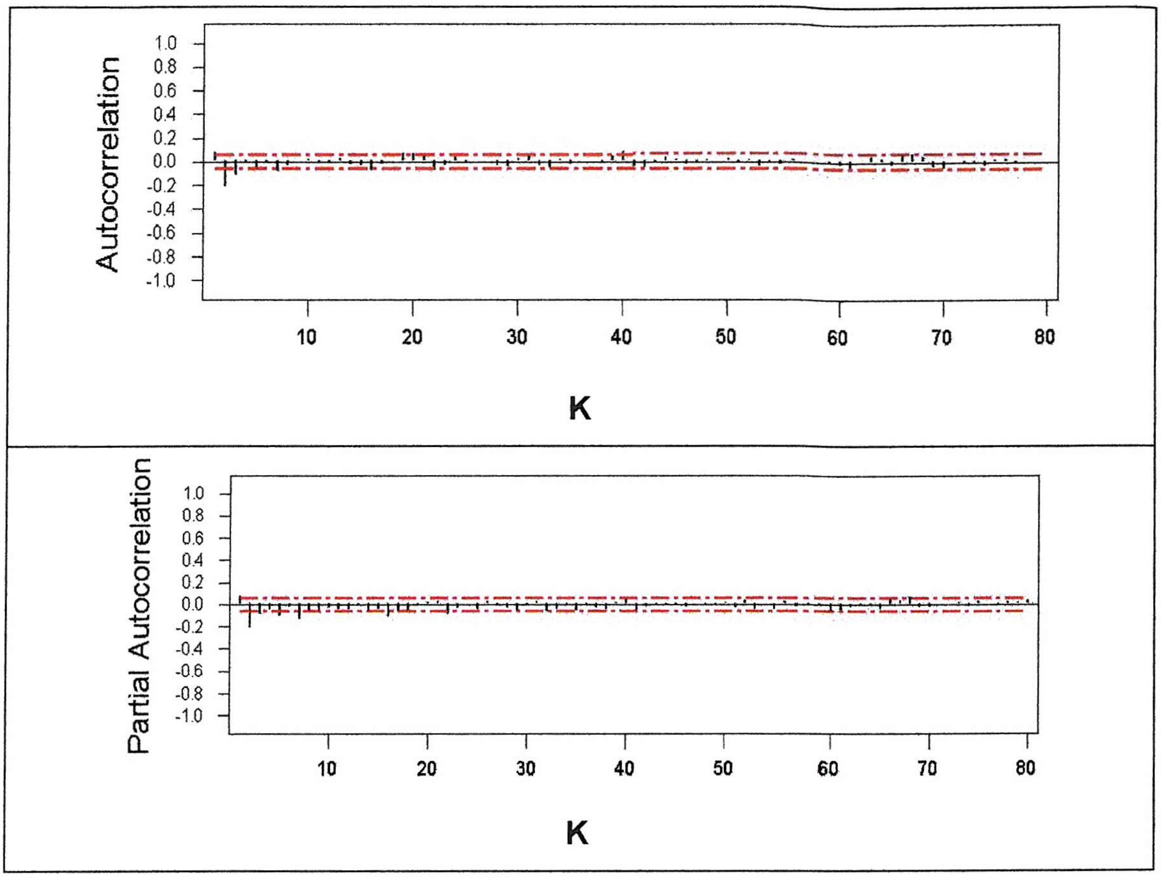

Figure 4-3 Autocorrelation and Partial Autocorrelation functions for the 1st difference of the >10MeV proton flux time series.

It would appear to a first approximation that the autocorrelation function becomes zero after three lags, whereas the partial autocorrelation function has non-zero terms beyond the 5th lag. Whilst this does not clearly represent any of the theoretical trends described in Table 4-1 it can be tentatively approximated by an ARIMA(0,1,3) model. (i.e. the autocorrelation cuts off after 3 lags whilst the partial autocorrelation tails off). An alternative interpretation is that both the autocorrelation and partial autocorrelation functions tail off, indicating a model with both auto-regressive and moving average terms. An ARIMA(3,1,3) model has therefore also been created.

ARIMA(0,1,3) and ARIMA(3,1,3) models were fitted to the sample data using MINITAB to perform least squares regression. The form of each model with their respective coefficients is summarised in Table 4-2. For completeness, and to act as a comparison, ARIMA models of form (012) and (011) were also constructed.

ARIMA(013)	$x_t = x_{t-1} + \mu + \varepsilon_t - \theta_1\varepsilon_{t-1} - \theta_2\varepsilon_{t-2} - \theta_3\varepsilon_{t-3}$	
	μ	-0.0002
	θ_1	-0.0344
	θ_2	0.2916
	θ_3	0.1999
ARIMA(313)	$x_t = x_{t-1} + \alpha_1 x_{t-1} + \alpha_2 x_{t-2} + \alpha_3 x_{t-3} + \varepsilon_t - \theta_1\varepsilon_{t-1} - \theta_2\varepsilon_{t-2} - \theta_3\varepsilon_{t-3}$	
	α_1	-0.531
	α_2	-0.018
	α_3	-0.499
	θ_1	-0.529
	θ_2	0.697
	θ_3	0.251
ARIMA(012)	$x_t = \mu + \varepsilon_t - \theta_1\varepsilon_{t-1} - \theta_2\varepsilon_{t-2}$	
	μ	-0.0002
	θ_1	-0.021
	θ_2	0.221
ARIMA(011)	$x_t = \mu + \varepsilon_t - \theta_1\varepsilon_{t-1}$	
	μ	-0.0002
	θ_1	-0.115

Table 4-2 Summary of ARIMA model form and coefficients after identification of possible models from examination of the sample autocorrelation functions.

4.3.3 ARIMA Model Performance

After fitting each model to the sample the NRMSE error was calculated. Results are shown in Table 4-3. Persistence has been used as a benchmark, and uses the current value of the proton flux as a forecast for the following value.

Model	NRMSE
Persistence	0.552
ARIMA(011)	0.549
ARIMA(012)	0.538
ARIMA(013)	0.530
ARIMA(313)	0.518

Table 4-3 Comparison of NRMSE for different ARIMA models and persistence calculated from the sample time series.

It can be seen that the sequential addition of MA terms to the model has given small improvements in accuracy, as has the addition of three AR terms. However, the NRMSE remains relatively high in all cases and indicates that none of the models actually provide an accurate forecast. The addition of three AR terms has resulted in only a 2% improvement in the NRMSE, suggesting that the form of ARIMA model cannot be significantly improved by adding further terms. Further terms will add additional complexity for very small gain in accuracy.

The best ARIMA model of those constructed is of (313) form with a NRMSE of 0.518, although this is only marginally superior to the NRMSE of persistence. Figure 4-4 shows the actual time series with the ARIMA(313) forecast plotted simultaneously for a period of approximately 100 days.

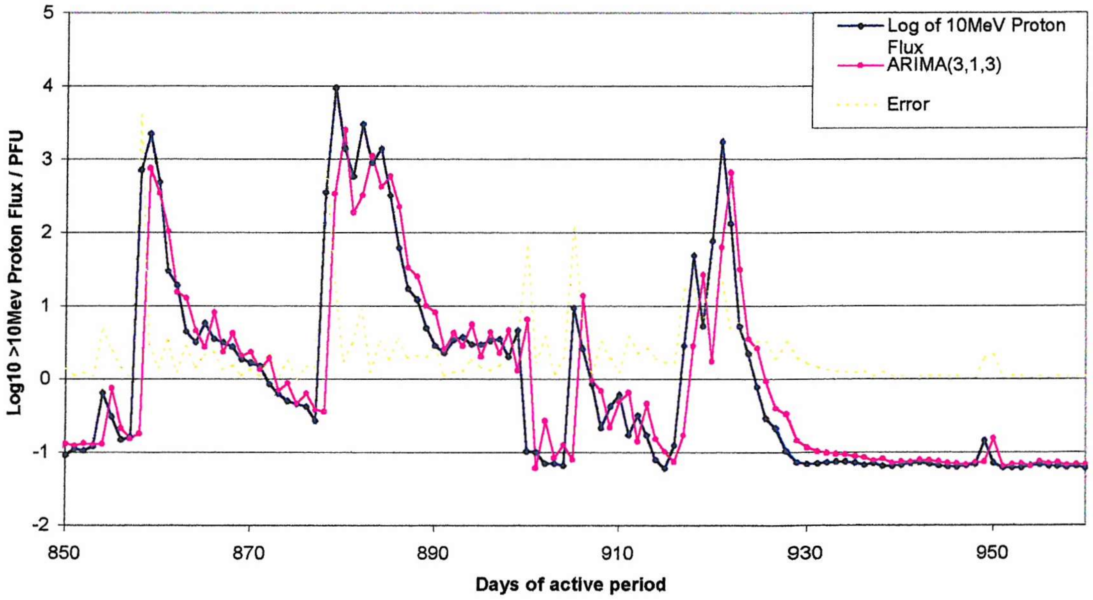

Figure 4-4 Predicted >10Mev Proton Flux and forecast >10MeV Proton flux using ARIMA(313) for a 125-day sequence of the time series.

Figure 4-4 shows that the resting proton flux (from around 930 days in the plot) is modelled well with almost zero error, but the large enhancements are not forecast at all. It is the sudden commencements of SPEs that dictate the high NRMSE of the model, and it can be seen that in each of the four major peaks (i.e. SPEs) the model lags the observed values, responding with a high forecast value only after high values have actually occurred and are used as inputs to the model. Success of the ARIMA technique is based on the principle that a time series is deterministic: i.e. that future values in the time series are related to current values of the time series. This appears to be the case when the proton flux is at a resting background, but is evidently not the case when an SPE enhancement occurs. Intuitively it is no surprise that the ARIMA model performs poorly in this respect as an inspection of the time series showed that the SPE enhancements were sudden, not trend-like, and appeared to have no dependence on the low background levels that preceded the enhancements.

The ARIMA(313) model has been extrapolated in order to provide forecasts for 2 to 4 days ahead by using the predicted times series as an input to the model. This is standard procedure for generating ARIMA forecasts for greater than 1 time step ahead. NRMSE errors for the 1-4 day forecasts are given in Table 4-4, and the predicted time series are plotted in Figure 4-5.

Forecast Lead	NRMSE
+1 Day	0.518
+2 Day	0.760
+3 Day	0.892
+4 Day	0.983

Table 4-4 NRMSE for ARIMA(313) model when extrapolating to higher lead times.

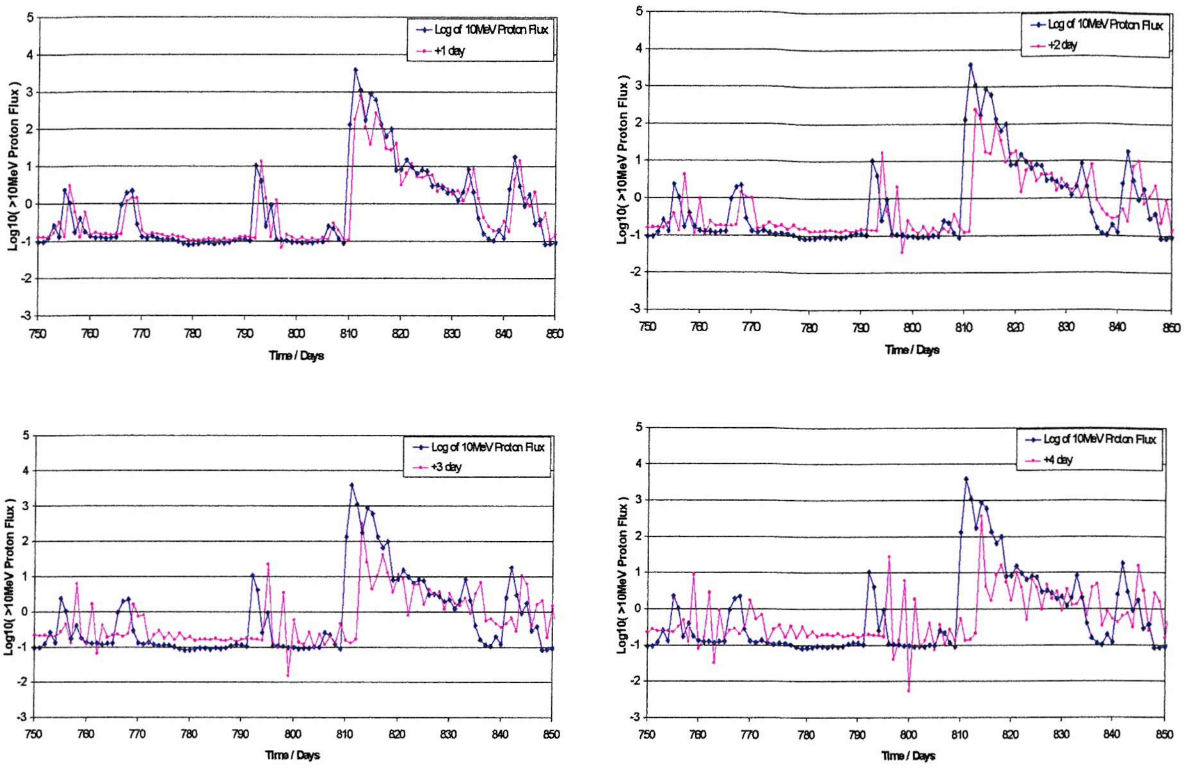

Figure 4-5 : +1, +2, +3, and +4 day ARIMA(313) time series forecasts for an arbitrary 100 day period plotted against the observed proton flux.

It can be seen that the forecasts rapidly break down when extrapolated to higher lead times. This is borne out in both the NRMSE, which approaches 1 for lead times of greater than two days, and the time series plot of the predicted flux, which tends to exhibit significant oscillation and deviation from the observed flux for lead times of greater than two days. It can be noted that none of the sudden enhancements in the observed proton flux are predicted by the model: there is always a lag.

The behaviour of the extrapolated predicted series shown in Figure 4-5 suggests that it is difficult to obtain lead times of more than a few steps ahead using ARIMA techniques, at least when applied to solar proton flux. The use of high resolution data, for example 1-hour averages, may improve the fit of ARIMA models as the higher sampling rate will result in a smoother transition between consecutive flux values during an SPE, although importantly, the initial enhancement itself is still likely to resemble a step increase, even on an hourly timescale. The problem though with using high resolution data is that the ARIMA technique can only predict 1-timestep ahead, and to generate longer lead times a series of extrapolations needs to be made on predicted values. The results from daily averages show that a large error is introduced and permeates quickly when extrapolating predictions, making forecasts of more than 2 or three time steps ahead very inaccurate. 1-hour averages are therefore not likely to be accurate when attempting to extrapolate to 24 time steps ahead or more (as would be required for a 1-day lead time).

4.4 Summary

The daily >10MeV proton flux time series over a 1200 day period commencing in May 1987 was found to be non-stationary, but could be made stationary by taking the 1st difference of the series.

Application of ARIMA time series forecast methods was applied to the differenced proton flux and is the first time that an ARIMA technique has been documented for the use of a solar proton flux time series. The technique did not result in an accurate forecasting model. ARIMA models were found to be only marginally better than persistence for a 1-day lead time, and had high NRMSEs of the order of 0.52 (for a 1-day lead time).

Increasing lead times by extrapolating forecasts on predicted data resulted in a significant loss in accuracy and NRMSEs approached 1 for predictions of greater than 2 time steps ahead. The use of higher resolution data may improve the fit of an ARIMA model, but it will be impossible to extrapolate to long lead times without introducing a significant error.

The high NRMSE was found to be dictated by large enhancements in the observed proton flux that were not predicted by the ARIMA models. The ARIMA model lags observations and only responds with high predicted values after an enhancement has occurred. The implication is that solar proton events are not predictable from a proton flux time series, meaning that the values of proton flux prior to SPEs are not indicative of the fact that an SPE is about to occur.

The ARIMA time series prediction technique fails to produce accurate results because it fails to predict the sudden onset of an SPE. Examination of the proton flux time series confirms that SPEs resemble significant step enhancements in a relatively invariant background flux level.

5. A CLASSIFICATION APPROACH TO SPE PREDICTION

This chapter revisits the concept of a classification prediction model in which it is intended to correlate the behaviour of solar variables with the occurrence of an SPE. Results from the previous chapter are drawn on in order to show that SPEs can be approximated as discrete phenomena, and a classification approach to the prediction problem is outlined in detail with reference to neural network models as classification tools.

5.1 Binary Representation of SPEs

The previous chapter found that SPEs were not predictable from a proton flux time series because their occurrence was sudden rather than trend-like. The nature of SPEs therefore suggests that they are better represented as discrete occurrences rather than as trends in a continuous flux-time profile. Whilst the flux and fluence of an SPE do vary significantly between different events, all SPEs are common in that their flux and fluence is significantly greater than background, and their commencement is often sudden, equivalent to a step impulse. It can be noted that both the PROTONS and Garcia models make a yes/no probability prediction as to whether an event will occur or not, further justifying the representation of SPEs as discrete phenomena.

Taking this a step further, the proton flux time series can be reduced to a binary series, composed of SPEs and non-SPEs. Figure 5-1 plots the proton flux time series for 1991 (chosen arbitrarily) in conjunction with a 'binary' time series, which has classified daily proton flux values of >1.0 p.f.u. as 'SPEs' and daily flux values of <1.0 p.f.u. as 'non-SPE' values. (A daily averaged flux level of >1.0 p.f.u. was used to identify SPEs in a previous listing of SPEs, which is why it has been used here for this illustrative example[*]). It can be seen that such a representation clearly identifies all the significant peaks in the proton flux time series.

---

[*] Dr. Stephen Gabriel, private communication.

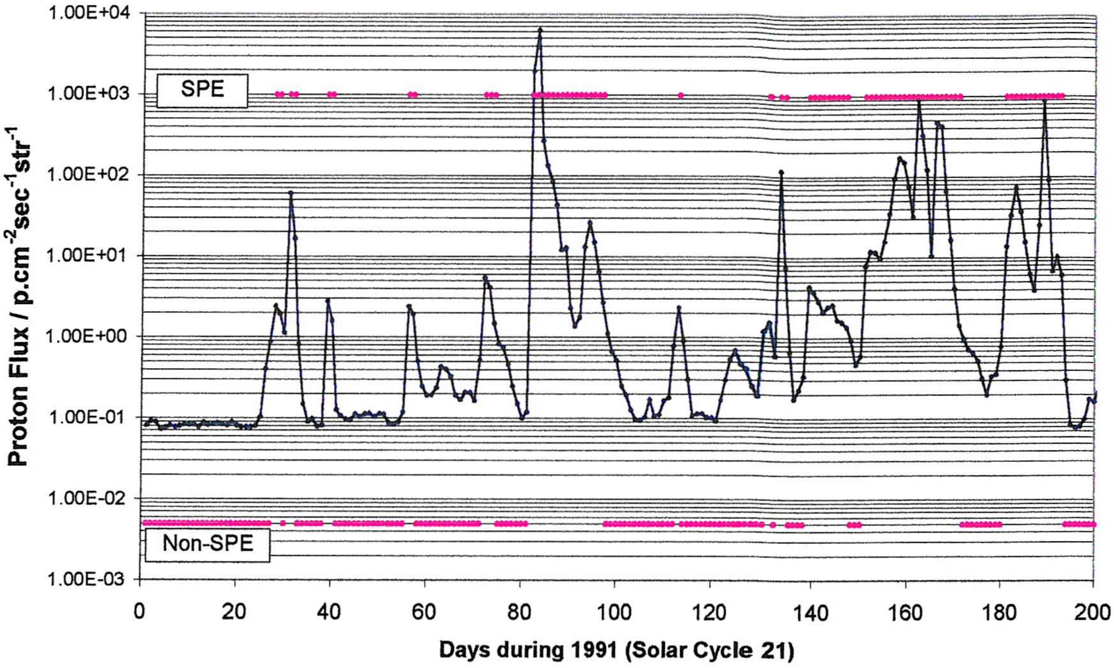

Figure 5-1 Daily proton flux time series during 1991 represented as a binary series of 'SPEs' and 'non-SPEs'.

Given that the primary aim is to predict the onset of an SPE (rather than a flux or fluence value) it does not matter that a binary time series does not represent specific flux values. The binary series simply identifies points in time at which the proton flux is deemed significantly high to warrant forecasting as an SPE occurrence.

5.2 The Classification Approach

The representation of SPEs as discrete events lends itself well to a classification approach to the forecasting problem, in which there are just two possible forecasts: that of an SPE or a non-SPE. The aim is to distinguish between these cases using input variables taken from before the time of the event. This is analogous to current flare forecasting models which predict the likelihood of an SPE from the characteristics of an x-ray flare.

In principle the technique will take inputs at time 't' over a given time window of length 'w'. The input window stretches back in time to time t-w, and is used to make a prediction for time t+T, where T is the lead time of the model. The technique can easily be applied to a rolling timeline by moving the input window forward in time by some increment, dt, to produce a prediction for time t+dt+T. This is summarised in Table 5-1 and shown graphically in Figure 5-2. Note that this technique always has a constant lead time, unlike current SPE prediction models, which either do not stipulate a lead time (Garcia model), or which have a variable lead time dependent on the inputs (PROTONS model).

Run Time	Input window span	Prediction time
$t = t_0$	$(t = t_0 - w) \longrightarrow (t = t_o)$	$t = t_0 + T$
$t = t_0 + \Delta t$	$(t = t_0 + \Delta t - w) \longrightarrow (t = t_o + \Delta t)$	$t = t_0 + \Delta t + T$
$t = t_0 + 2\Delta t$	$(t = t_0 + 2\Delta t - w) \longrightarrow (t = t_o + 2\Delta t)$	$t = t_0 + 2\Delta t + T$

Table 5-1 : Principle of the classification technique showing the time span of the input window relative to the time for which a prediction is made, applied to a rolling timeline.

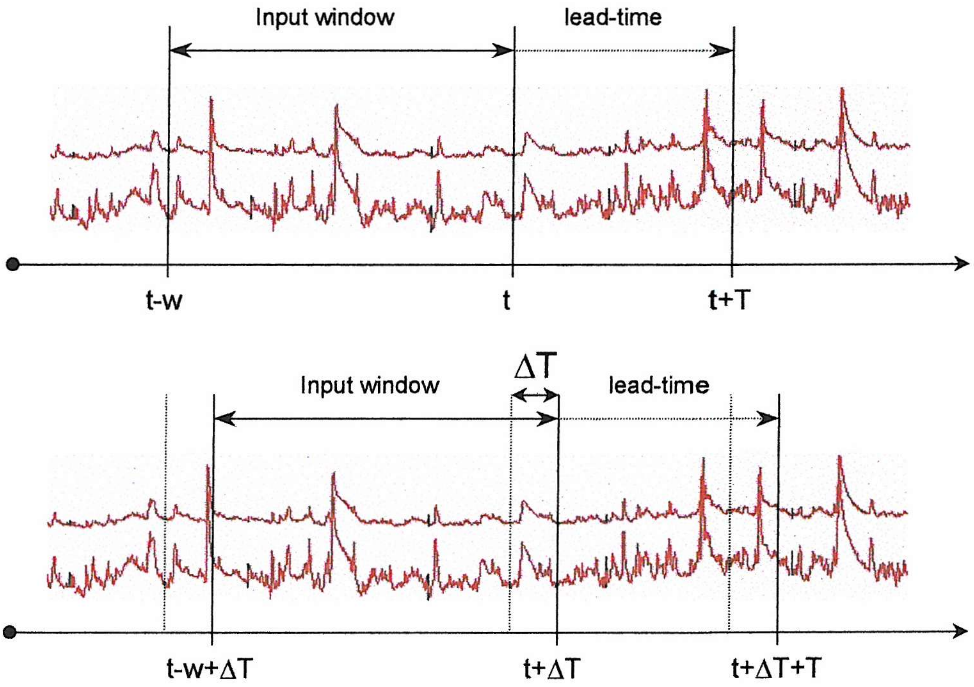

Figure 5-2 Graphical illustration of the proposed classification forecast model applied to an arbitrary time series (in this case solar x-ray flux) at time 't' and at time 't+dt', where 'dt' is the step time of the model. A binary 'SPE'/'Non-SPE' forecast is issued for time t+T.

The approach benefits from a relatively simple problem structure in that the possible outputs are restricted to just two categories. Success of the technique will of course depend on the existence of some behaviour or precursor in the input variables that allows the cases of 'SPE' and 'non-SPE' to be differentiated between.

The construction of such a classification model must be based on historical data, and will require the collection of examples for each of the possible outcome categories. Specifically this means that a list of SPEs and their occurrence times will be needed, and correspondingly, a list of times at which no SPE occurred is required. Solar quantities (i.e. potential predictor-variables) must then be collected at these times (or more exactly, just prior to these times), to create a dataset of input windows pertaining to the classes of 'SPE' and 'non-SPE'. The aim is that a comparison of the predictor variables during the input windows will show a difference in behaviour and allow the two cases of 'SPE' and 'non-SPE' to be differentiated between.

5.2.1 Input Variables

Clearly, the input variables to the classification model need to have the potential to contain behaviour related to an SPE, but there are also significant practical issues to consider. The dataset on which the model is constructed will be based on as many SPEs as possible, and hence will span several decades. As consistency and uniformity of the predictor variables is highly important this implies that candidate variables should also have a long monitoring history, ideally of the order of decades. The intention to develop a real-time forecasting model also means that in order to be operational, any input variables that are used must be available in real-time or as near to as possible.

An obvious choice of input variable is GOES x-ray data. GOES x-ray sensors have been operational since 1974, and from 1986 onwards data is available in an off-the-shelf product. Recent measurements are available on-line from the SEC with an approximate delay of 10 minutes from real-time. As well as satisfying practical requirements, solar x-ray data has the obvious flare relation to SPEs (although in the classification model the input window would end well before any classical 'proton producing' flare), and in addition the ratio of the two GOES x-ray channels is a measure of the spectral hardness of solar emissions, which has been strongly associated with SPE occurrence [*Kiplinger*, 1995].

Another candidate for use as an input variable is the solar radio flux, which has been shown to reach a maxima on the day of SPE occurrence and therefore may have potential for forecasting SPEs on a timescale of days [*Das, Chakraborty, and Dasgupta*, 1987]. Solar radio fluxes have been monitored routinely by ground based observatories since the 1940s and archived daily averages exist from this time, hence there is little problem in obtaining values for a large number of SPEs. Being of minimum 1-day resolution solar radio flux values are published for the preceding calendar date (effectively a 1-day lag from the current time), which is not ideal for real-time operation, but the potential link between solar radio flux activity and SPEs still means that they are worth considering as an input variable.

5.3 Neural Networks as a Classification Technique

Neural networks provide a novel and potentially powerful solution to the classification problem that has been set-up in order to try and forecast SPEs. Their constituent component – the node – is extremely simple to model, yet by combining a large number of these into an organised structure, the resultant network is capable of solving relatively complex problems.

Fundamentally, neural networks differ from conventional modelling techniques in that their solutions are learnt and not programmed. Knowledge is gained by presenting a network with existing examples, and via a learning algorithm the free parameters within the network are adjusted such that specific input vectors are mapped onto the desired responses. Such an approach means that no detailed statistical model is made for the input data: it is the raw dataset itself that is used to derive the network's parameters. Typically, neural networks are often applied to problems in which the inter dependency of variables is unknown and when there is little a-priori knowledge of a system.

A principle advantage in the use of neural networks is that they are capable of producing non-linear solutions, and this may have particular benefits when the inputs to a model are known to originate from non-linear mechanisms (as is probably the case in solar particle phenomenon). Classification problems in which classes are not linearly separable cannot be solved by linear techniques, but can be solved by neural models [*Haykin*, 1999]. This ability of neural networks has been borne out in classification models for remote sensing, in which neural models have been shown to outperform linear regression techniques for certain datasets [*Cipollini*, 2000], [*Paola*, 1995].

Another advantage of neural networks is their ability to easily incorporate inputs from different sources. The use of predictor variables which have very different orders of magnitude can cause errors when calculating the coefficients in linear regression models and give uneven weighting of parameters in the model. By contrast, inputs to a neural model are scaled to within the same limits by the input layer of nodes, allowing many inputs to be incorporated easily, regardless of their origin or data-type (i.e. numerical value or discrete class-value).

The use of neural networks in conjunction with solar activity is not new, although it is a recent area of development. Gothoskar, describes an artificial intelligence technique for the detection of interplanetary disturbances, in which neural networks are trained to recognise disturbed power spectra in radio scintillations [*Gothoskar and Khobragade*, 1995], and in the late 1990s Costello developed a model (which is currently operational on SEC web-space⊕) to run in real-time using a neural algorithm to predict the Kp index from solar wind parameters [*Costello*, 2001]. A similar model is presented by Gleisner et al which uses RBF networks to predict local disturbances in the geomagnetic field from solar wind velocity [*Gleisner and Lundstedt*, 2001]. In very recent work, two independent studies have used ANNs to predict DST indices, reporting smaller errors than with existing models [*Lundstedt, Gleisner, and Wintoft*, 2002] [*Watanabe, Sagawa, Ohtaka, and Shimazu*, 2002], and a neural pattern recognition algorithm has been used to identify solar flares from full disk images in order to act as an unsupervised flare patrol system [*Borda, Mininni, Mandrini, Gomez, Bauer, and Rovira*, 2002].

Given that neural techniques have already been successfully applied to the prediction of space weather variables, their use in the prediction of SPEs is an obvious progression. This is especially true taking into account the fact that the problem has been set-up as one of classification (an area in which neural networks have been shown to excel) and that there is no a-priori knowledge of the relation between the predictor variables and SPEs.

⊕ http://www.sec.noaa.gov/rpc/costello/

5.3.1 MLP Networks

Neural networks are composed of perceptrons, which constitute artificial models of biological neurons. The concept of the perceptron was proposed by Rosenblatt and forms the basis for neural network theory [*Rosenblatt*, 1958].

Each perceptron is composed of a set of input links, a summing junction and an activation function. Each input link has an associated weight such that the signal xj at the input of link *j* connected to perceptron *k* is multiplied by weight w_{kj}. The summing junction adds each of the input×weight products and passes the result to the activation function. The result of the activation function is then passed out of the perceptron as an output. The arrangement is summarised in Figure 5-3.

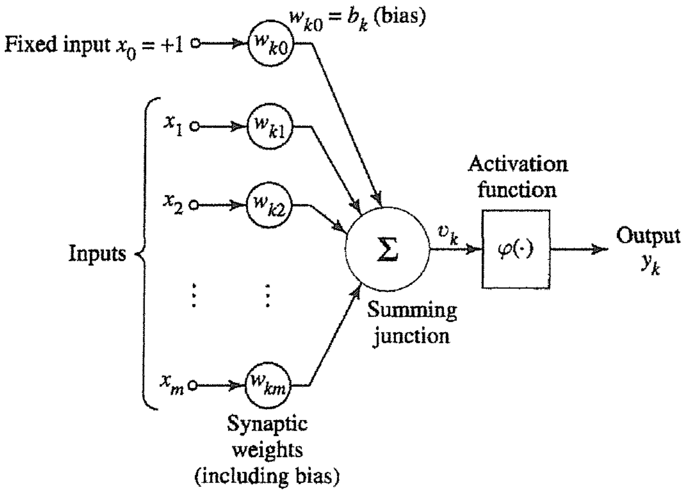

Figure 5-3 Representation of the Perceptron

The activation function defines the perceptron output as a function of the sum of the perceptron inputs, and whilst it can be a linear function, the non-linearity of a neural network is achieved by using a non-linear activation function. The sigmoid function is commonly used for this purpose, and is defined as:

$$\varphi(v) = \frac{1}{1 + \exp(-av)}$$

in which a is a slope parameter defining the rate of change of φ(v) with respect to (v), where v is the sum of the perceptron inputs. A graph of the sigmoid function is given in Figure 5-4.

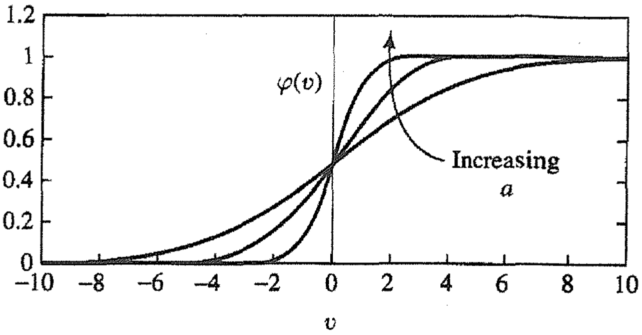

Figure 5-4 The sigmoid function.

By convention, each perceptron is modelled as always having one of its inputs fixed at 1 with weight w0. This serves as a bias, which has the effect of shifting the activation function along the x-axis, such that the threshold is non-zero. The effect of the activation function is to map the sum of the perceptrons' inputs to an output value of between 0 and 1. In the case of the sigmoid function the mapping is a continuous function, and is thus differentiable: a property which allows use of the back propagation learning algorithm in the Multi Layer Perceptron network (as shown below).

The Multi Layer Perceptron (MLP) class of neural networks is composed of an input layer of source nodes, an output layer of perceptrons and any number of hidden layers of perceptrons between the two. Source nodes serve to receive external inputs to the network and transmit them to the first hidden layer: they have no computational value. Outputs from the first hidden layer are then passed to the next hidden layer and so on. The output from the output-layer constitutes the overall response of the network to the input signal. The addition of hidden layers of perceptrons increases the number of dimensions of node interactions, essentially adding degrees of freedom to the network enabling higher order statistics to be extracted from the inputs. An example of an MLP structure is given in Figure 5-5.

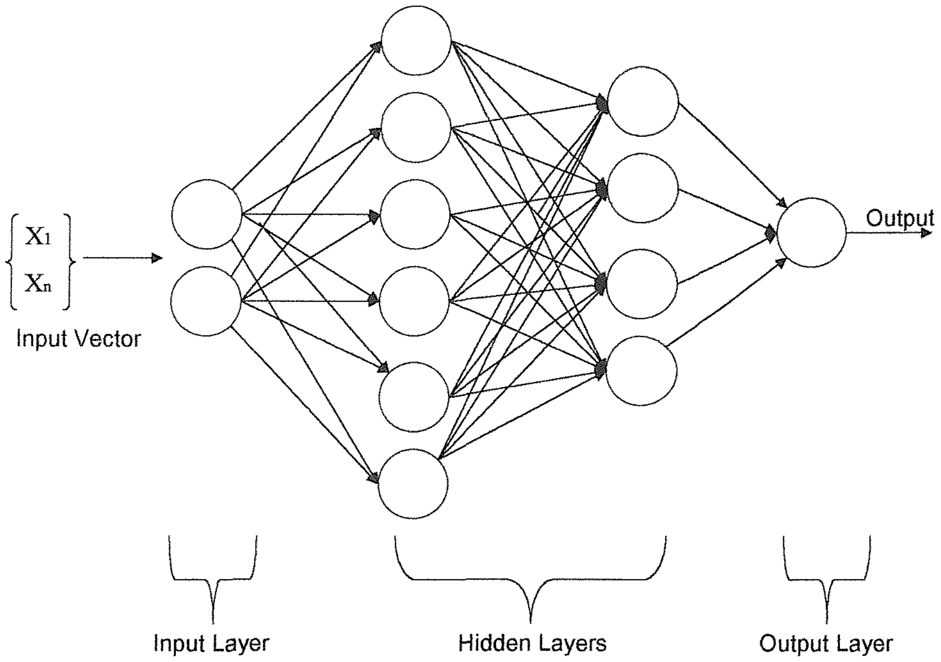

Figure 5-5 Example of MLP architecture

The MLP network adapts its free parameters, i.e. its weights, to learn a solution via the principle of error correction learning. Given that an MLP network receives an input vector $\mathbf{x}(n)$ at iteration n, it will produce an output vector equal to $\mathbf{y}(n)$. Comparing the output vector to the desired response $d(n)$ produces an error signal, $e(n)$ where:

$$e(n) = d(n) - y(n)$$

The weights of a network must be altered so as to reduce the error of the network output. Figure 5-6 shows the variation of the squared error with weight values for a simple 2-weight system. Before training commences the weights are set at small random values placing the system at any point on the error surface. During training the network seeks to adapt its weights so as to follow the steepest descent of the error surface.

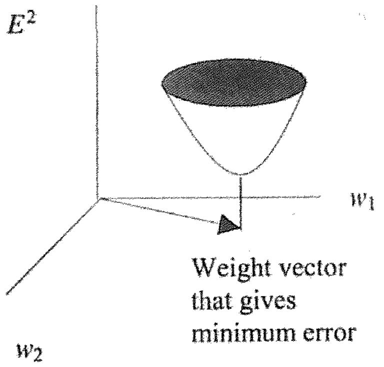

Figure 5-6 Visualisation of error surface for a 2-weight system

The learning process is implemented computationally via the back propagation algorithm. The process is a supervised learning technique that requires a training dataset, such that each input vector has an associated target output that can be compared against the actual output of the network. During the training process each input vector within the training dataset is presented at random to the network and an output is produced by propagating the input signal forwards through the network. An error signal is then derived by comparing the network output to the target response, and this signal is propagated backwards through the network, permitting the network weights to be altered by a small amount such that the error signal is minimised on following forward passes. The presentation of training examples continues (cycling through the training set) until values for the weights stabilise and a steady state is reached.

The back propagation algorithm is stated without proof as:

$$\Delta w_{ij}(n+1) = \eta(\partial_j o_i) + \alpha\Delta w_{ij}(n) \qquad \text{5-1}$$

Where $\Delta w_{ij}(n+1)$ represents the change in weight ij for pattern n+1. η is a learning coefficient which determines the step size used to traverse the error surface. δ_j is defined as $-\partial E/\partial net_j$ and is the rate of change in error (E) with respect to the sum of the inputs for node j (net_j). o_i is the output from node i and $\alpha\Delta w_{ij}(n)$ is a momentum term that adds a proportion of the previous weight change (the proportion being dependent on the value of α). The weight change for the pattern n+1 is thus dependent on the previous weight change for pattern n and prevents oscillation of the weight changes due to local error minima.

A derivation and detailed implementation of the back propagation algorithm can be found in [*Haykin*, 1999]. The algorithm is a generalised form of the Widrow-Hoff rule [*Widrow and Hoff*, 1960].

5.3.2 Radial Basis Function Networks

The Radial Basis Function (RBF) approach works on the same principle as an MLP network, but treats the problem as one of curve fitting in multi dimensional space. Input vectors to an RBF network are projected via non-linear functions into a high-order multi-dimensional space, the principle being that once a pattern classification problem is cast into a high-dimensional space it is more likely to have a linearly separable solution.

An RBF network usually has only three layers: an input layer of nodes which receive the input vector, a hidden layer, in which each node consists of a function that acts as the basis for the projection of the input vector into hidden space, and an output layer, in which the resulting points in the hidden-space are mapped back onto meaningful values.

In the case of a binary classification problem an RBF network performs a mapping from an m_0 dimensional input space to a one dimensional output space such that:

$$\Re^{m_0} \longrightarrow \Re^{1}$$

The mapping takes the form of a hyper-surface, Γ in which the output is given as a multi dimensional function of the input vector. The aim, via the fitting of a curve to a set of training data, is to find the function which approximates this surface. Specifically, in the RBF technique, the function takes the form of a linear sum of functions given by:

$$F(\mathbf{x}) = \sum_{i=1}^{m_1} w_i \varphi(\|\mathbf{x} - \mathbf{c}_i\|) \qquad \textbf{5-2}$$

Where $\varphi(..)$ is a set of m_1 arbitrary functions, denoted radial-basis functions, positioned at centres $\mathbf{C}_i$. (Typical radial basis functions might be a cubic function $\varphi(r) = r^3$,or the

Cauchy function $\varphi(r) = (1+r^2)^{-1}$). $\mathbf{w}_i$ are coefficients chosen so as to minimise the error *E*:

$$E = \sum_{i=1}^{N} \left(F(\mathbf{x}_i) - \mathbf{d}_i\right)^2 \qquad \textbf{5-3}$$

Where $\boldsymbol{d}_i$ is the desired response vector from input vector $\mathbf{x}_i$. For a training set of N examples, consisting of N input vectors $\mathbf{x}_i$, and N target responses $\mathbf{d}_i$, values for the coefficients, w_i, can be found by solving the resulting set of linear equations:

$$\begin{bmatrix} \varphi_{11} & \varphi_{12} & \cdots & \varphi_{1m} \\ \varphi_{21} & \varphi_{22} & \cdots & \varphi_{2m} \\ \vdots & \vdots & \vdots & \vdots \\ \varphi_{N1} & \varphi_{N2} & \cdots & \varphi_{Nm} \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_m \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \\ \vdots \\ d_N \end{bmatrix} \qquad \textbf{5-4}$$

$$\varphi \mathbf{w} = \mathbf{d} \qquad \textbf{5-5}$$

$$\mathbf{w} = \varphi^{-1} \mathbf{d} \qquad \textbf{5-6}$$

The complexity of the surface function is controlled by the number of centres. At maximum the number of centres can equal the number of input vectors in the training set, resulting in a highly complex hyper–plane which passes through every point in the dataset, but this is rarely desirable as it leads to any noise in the dataset also being modelled. The optimal number of centres can be found via experimentation for any given problem. The complexity and inherent form of the hyper surface can also be altered via the use of different types of RBF functions.

The locations for the centres functions can be determined at random, but superior generalisation can result if the centres functions' locations in multi dimensional space are chosen based on properties of the training set. The k-means clustering algorithm provides a method for achieving this, and results in centres being placed only at points in multidimensional space where significant data are present.

The k-means algorithm proceeds by drawing an input vector from the training set and computing the Euclidean distance in multidimensional space between this point and all available centres in order to find the closest centre. The position of this centre is then adjusted by a small amount so as to bring it closer to the point of the input vector. The iteration continues until there is negligible change in the centre positions, the result of the algorithm being that the centres drift towards regions in multidimensional space where there is a large density of data points. The RBF functions are therefore placed in regions where it is important to model the hyper-plane more accurately.

5.3.3 Linear Regression

Linear regression represents a classical-statistics alternative to the problem of SPE prediction via classification. Linear regression is typically applied within the scope of experimental measurements in which it is desired to find a function relating variable x to output y. Given a scatter plot of variable y as a function of x, regression methods can be used to derive the equation for the line of best fit equating y as a function of x. The technique can be extrapolated to accommodate any number of input variables in which it is desired to know 'y' as a function of x_1, x_2, ..., x_p and is termed multiple linear regression.

The method of multiple linear regression obtains coefficients for the linear equation defining a criterion variable 'y' as a function of 'P' predictor variables (X_1, X_2, ..., X_p) such that:

$$Y = b_0 + b_1X_1 + b_2X_2 + b_3X_3 + \ldots\ldots + b_pX_p \qquad \textbf{5-7}$$

The coefficients b_0, b_1,.... ,b_p are derived from a dataset, (equivalent to a training dataset), of 'n' observations, where each observation consists of 'p' predictor variables and one criterion variable. The coefficents are found by minimising the squared error of the function. The objective function 'F' can thus be defined as:

$$F = \text{minimise}\sum_{i=1}^{n} e_i^2 = \min\sum\left(b_0 + \sum_{j=1}^{p} b_j X_{ij} - y_i\right)^2 \qquad \textbf{5-8}$$

where i indicates the observation and j the specific predictor variable. The equation is solved by taking the (p+1) derivatives of the objective function with respect to all the b_j's, setting them to zero and solving for the unknowns. The result is a set of (p+1) simultaneous equations which can be solved for all b_j (j = 0, 1,... p).

5.3.4 Principal Components Analysis

The technique of principal components analysis (PCA) can be used as a data pre-processing tool in order to produce a dataset that is more receptive to a neural or linear classification technique. The process of a principal components analysis identifies the primary constructs within a dataset, allowing noise to be removed and the total number of dimensions within a dataset to be reduced without significant loss of information.

Given an untransformed dataset matrix **X** of N rows (examples) and m dimensions (columns) the corresponding correlation matrix of **X** is given by:

$$\mathbf{R} = \mathbf{X}\mathbf{X}^{\mathrm{T}}$$

In the ideal case **R** will consist of a diagonal matrix indicating that there is zero inter-correlation between the predictor variables of the dataset. In practice this is rarely the case and some non-diagonal elements may be significantly non-zero, indicating a high inter-correlation between variables. This can lead to irrational coefficients when building models, poor performance, and an overly numerous number of input variables (i.e. overly complex models).

The aim of a PCA is to create a new dataset **A** with a diagonal correlation matrix by projecting the original dataset **X** onto a set of axes, termed principal components. Matrix **A** is therefore determined by having minimum inter-correlation between predictor variables (and thus maximum variance contained within separate predictor variables).

Given the pre-requisite to find a diagonal correlation matrix of **A**, Ψ, the aim is to find a matrix **Q** such that:

$$\Psi = \mathbf{Q}'\mathbf{R}\mathbf{Q}$$

where **R** is the correlation matrix of the original dataset. A unique solution is given subject to the constraint that that the vectors of **Q** are orthogonal to each other and leads to the expression:

$$\mathbf{RQ} = \mathbf{Q\Psi}$$

where **Q** is composed of the eigenvectors of **R** ($[\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_m]$) , and Ψ is a diagonal matrix with the squares of the eigenvalues as elements of the principal diagonal. A full derivation of the principal components solution is given by Haykin [*Haykin*, 1999].

The projected dataset matrix **A** is then given by:

$$\mathbf{A} = \mathbf{Q}^{\mathrm{T}}\mathbf{X} = \mathbf{X}^{\mathrm{T}}\mathbf{Q}$$

with each row of **A**, (**a**) being related to each row of **X**, (**x**), by:

$$\mathbf{a} = \left\lfloor \mathbf{x}^{\mathrm{T}} q_1, \mathbf{x}^{\mathrm{T}} q_2, \ldots\ldots, \mathbf{x}^{\mathrm{T}} q_m \right\rfloor$$

where **X** is a dataset of m dimensions (columns). A dataset of new 'm' dimensional vectors is thus created by projecting each row of the original dataset using the eigenvectors calculated from **R** (the correlation matrix of the original dataset **X**).

Each element, i.e. dimension, of the new vector, **a**, is a principal component, the *j*th principal component being given by:

$$\mathbf{a}_j = \mathbf{x}^{\mathrm{T}}\mathbf{q}_j \qquad j = 1, 2, \ldots m$$

Each principal component is thus composed of a contribution from every one of the original dimensions of vector **x**, the weight of contribution being determined by the values of the elements within eigenvector $\mathbf{q}_j$.

Given a vector **x** of 'm' dimensions (i.e. of length m), it can be reconstructed from the principal components as:

$$\mathbf{x} = \mathbf{Qa} = \sum_{j=1}^{m} a_j \mathbf{q}_j$$

However, an approximation to **x** can be obtained by truncating the expansion to:

$$\hat{\mathbf{x}} = \mathbf{Qa} = \sum_{j=1}^{l} a_j \mathbf{q}_j$$

where $l<m$. The original vector **x** of length 'm' can thus be represented by a transformed vector of length 'l' (i.e. a vector of fewer dimensions).

The amount of variation expressed by each eigenvector $\boldsymbol{q}$ is measured by the value of its corresponding eigenvalue (i.e. the values within the matrix Ψ), thus the dominant eigenvalues correspond to the principal components that represent the majority of variation in the untransformed dataset. A significant amount of the variation is usually accounted for by only a small number of the total eigenvectors, thus the variation within the initial dataset can be expressed relatively accurately by a transformed dataset of far fewer dimensions, (each dimension being one of the significant principal components). Eigenvectors with low eigenvalues simply represent noise within the dataset.

Performing a PCA allows a dimension-reduction of the number of inputs within each input vector, whilst at the same time, increasing the signal to noise ratio of any (potentially) useful variation.

6. DATASET

This chapter describes the construction of a dataset on which the classification model described previously can be based. Principally it defines a list of SPE occurrences and a list of times at which the proton flux is at a background level. Solar x-ray and radio flux measurements have then been gathered at these times to provide predictor data relating to the two cases.

6.1 Solar Proton Data

Solar protons have been monitored by space borne instruments since 1963 and as a result there is a large amount of historical data available from which to source events from. Space borne solar proton data is available from two prime sources:

6.1.1 The IMP Spacecraft

The Interplanetary Monitoring Platform (IMP) series of satellites have been operational since 1963 and are equipped with a solar proton monitor capable of monitoring protons in the >1, >2, >4, >10, >30 and >60 MeV energy bands and above. The IMP satellites have been operational in a highly elliptic earth orbit with a 12 day period and apogee and perigee at ~45 and ~25 earth radii respectively. The orbit intentionally enters and exits the earth's magnetosheath and magnetotail meaning that measurements can be taken directly in the solar wind outside the influence of the earth's magnetic field. The energetic particle data from the IMP satellites is taken in the solar wind part and the non solar wind part of the orbit, although the dataset itself makes no distinction as to where the satellite was when measurements were made. IMP satellite data can be retrieved and plotted free of charge from the OMNI-web data site facility operated by NASA. Proton data is available in a minimum resolution of 1-hour averages.

6.1.2 The GOES Spacecraft

The SMS/GOES (Geosynchronous Operational Environmental Satellites) series of spacecraft occupy a geo-stationary orbit at an altitude of 6.67 earth radii and are located between 75° and 135° West longitude. The main payload is an earth facing visual and infrared imager for terrestrial weather forecasts, but in addition the satellites carry a Space Environment Monitor (SEM) instrument package equipped to monitor solar x-rays, solar particles (alpha particles protons and electrons) and the earth's magnetic field. The SEM has provided data continuously since July 1974, although off-the-shelf data is available only from 1986 onwards. The Energetic Particle Sensor measures proton flux in numerous energy bands ranging from >0.6 to >685 MeV. The X-ray sensor (XRS) consists of an ion chamber detector providing whole-sun x-ray fluxes for the 0.5-to-4 Å (XS channel) and 1-to-8 Å (XL channel) wavelength bands.

6.1.3 Solar Proton Data Retrieval

IMP proton data was sourced directly from the OMNIWEB data site from which a minimal resolution of 1-hour averages could be viewed and downloaded.

GOES proton data was purchased on CD from the NGDC. Data from 1986 onwards was available as an off the shelf data product with a minimal resolution of 5-minute averages. In addition to this, archived GOES data was purchased covering the initial operation period of the GOES satellites between 1974–1986. Archived data from this period was not in a ready to use format and needed extensive processing before it could be used. This consisted of a binary-to-ASCII conversion, followed by a de-multiplexing routine to recover 3-second resolution values from the raw data stream. 5-minute and higher order averages were then calculated from the 3-second data. The procedures were performed via custom written C-code.

6.2 A Proton Event Listing

Data from the IMP and GOES satellites has been taken and used to construct a list of SPEs over the period spanning from 1965 to 1999. The task was facilitated by making use of a pre-existing list of SPEs derived from IMP satellite data spanning from 1965 - 1989, in which a calendar day was said to contain a proton event if the >10 MeV daily averaged flux exceeded 1.0 pfu[*]. A proton event was defined as continuing until the daily averaged flux fell below this threshold and remained there for 1-day. This listing was obtained from Dr. Stephen Gabriel and originates from an earlier study of proton fluences on which the JPL-91 model was partly based.

---

[*] pfu = proton flux unit = protons/cm^2 s sr

For the period spanning 1986 – 1999 GOES proton data were found to be of superior quality and resolution to IMP proton data. For this reason GOES proton data was used to derive a list spanning the period from 1986 to 1999. SPEs were identified by plotting the GOES data and recording times at which the >10 MeV proton flux was greater than 1.0 pfu.

Amalgamating the two lists resulted in an SPE listing which covered the period 1965-1999, encompassing 284 events over two and a half solar cycles. During the period 1986-1989 in which the lists overlapped, the lists were compared closely in order that no event was duplicated.

It is important to note that the IMP definition for an SPE is different to that which the SEC uses for its forecasting services, and reflects the fact that the JPL-91 model was primarily interested in modelling the proton *fluence* contribution from SPEs. The NGDC defines an SPE as having occurred if the >10MeV integrated proton *flux* is greater than 10 pfu for 3 consecutive 5-minute readings. The NGDC definition thus has a higher flux threshold by 1 order of magnitude but does not consider the fluence of an event. In the majority of cases an SPE will fall into both definitions, but the IMP definition is the most conservative in terms of flux. In order to maximise the size of the dataset, and to take advantage of the existing list, the IMP definition was preferred. The disadvantage is that small SPEs may pollute a dataset in which it is only important to consider large SPEs. This was addressed at a later stage by providing a way to filter out small SPEs.

6.2.1 Start Times of Proton Events

For the purposes of real-time prediction it was deemed important to pinpoint the start of a proton-flux enhancement in order to correlate accurately with other solar variables. The existing IMP list only determined proton events to their calendar day of occurrence, and it was felt that a higher degree of accuracy was required. The start time of an event was thus defined as:

"The two hour period within which the >10MeV integral proton flux associated with a proton event begins to rise above background".

The different orbits occupied by the GOES and IMP spacecraft mean that an SPE can be monitored at one platform before it is seen at the other. It was therefore necessary to standardise proton data to one platform, and given that GOES x-ray measurements were to be used as a predictor the start times for the SPEs were all found in GOES proton data.

The integral >10MeV proton flux was plotted from GOES data archives for all proton events occurring after 1974 (i.e. even those events defined from the IMP listing). This ensured that the start times were all based on data from the GOES platform and would be co-ordinated with x-ray measurements from the GOES x-ray monitors.

For proton events that could be plotted from both GOES data and IMP data a comparison was made between event start times, and it was found that in over 75% of cases the difference between these start times was less than 2 hours. Thus, in the few instances where start times could not be derived from GOES data, a start time was derived by plotting IMP data from the OMNIWEB data site. Prior to 1974 no GOES data is available hence start times were not found for any SPEs which occurred prior to this.

6.2.2 Classification of Proton Events

In an effort to sustain a good quality dataset each of the proton events were classified as belonging to one of five classes according to the character of their flux time curves. These classes were purely qualitative but were created to serve as a method of filtering out any proton events that were small or badly defined. The five classes are described below:

Class 1 – Suitably Well Defined

The Event has a distinct start and is easily distinguishable from background levels

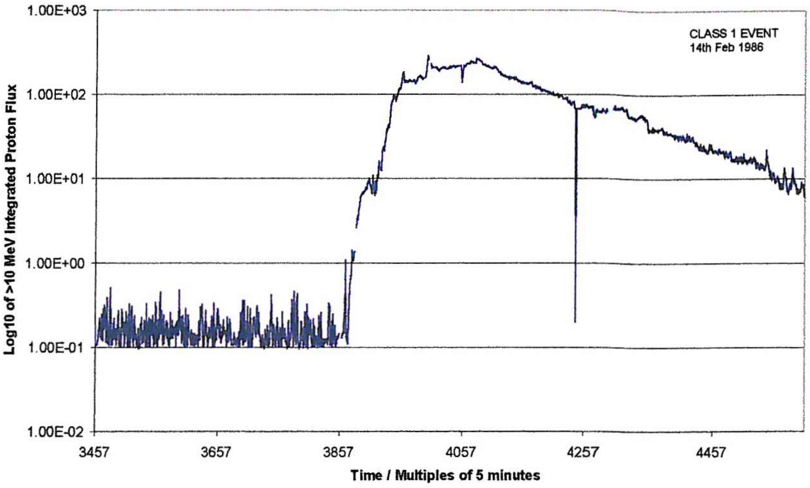

Class 2 – Gradual Rise

The event has a gradual rise to its peak flux spanning several hours and is susceptible to error in the definition of its start time.

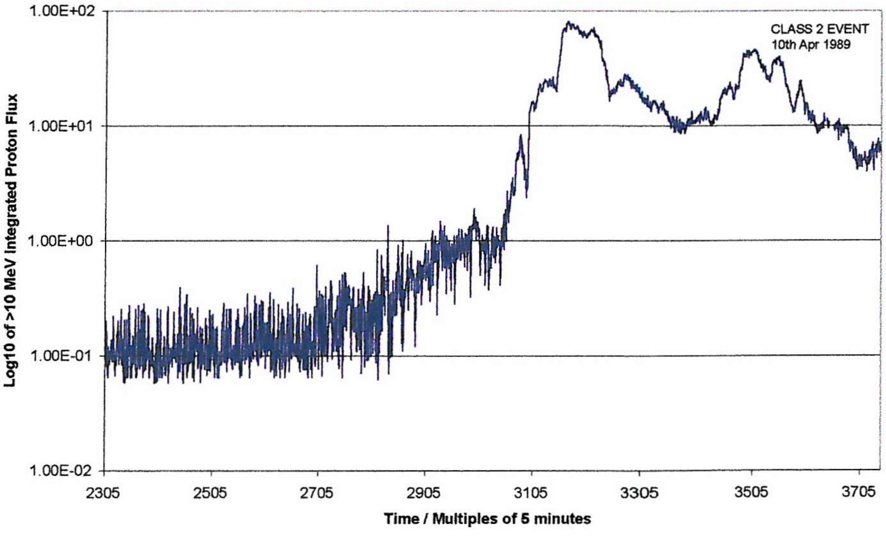

Class 3 – Not Well Defined

The proton flux from en event is below the SEC threshold of 10 pfu, or data quality is poor and renders the start time of the event difficult to pinpoint.

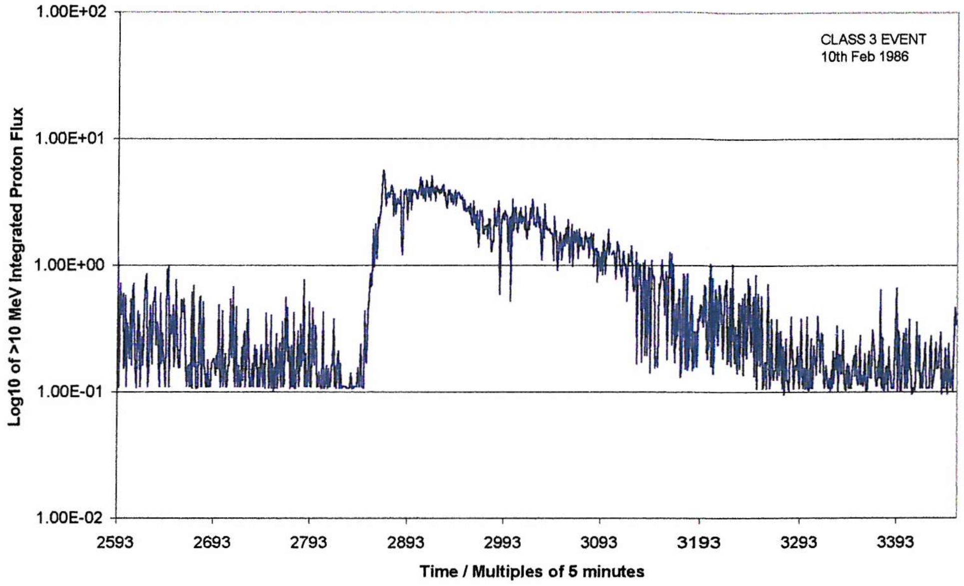

Class 4 – In Tail

The proton event occurs in the tail of a previous event and is not an isolated enhancement.

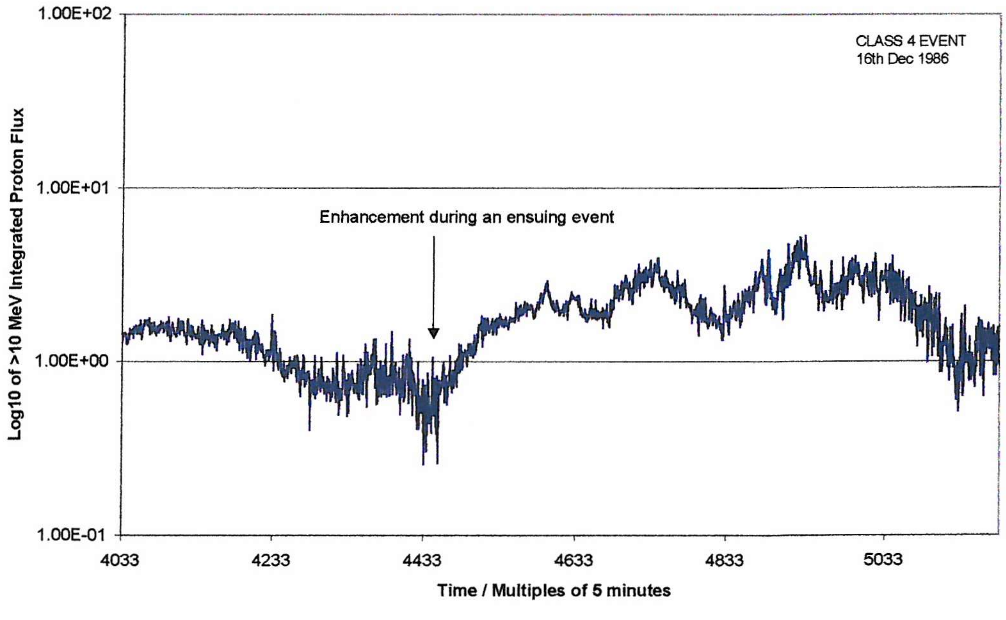

Class 5 – Unsuitable

The event is non-existent or bad/missing data spans the start time of the event.

As mentioned above, the definition of proton events by the IMP listing has a threshold one order of magnitude below that of the SPE definition used by the SEC. This means that the current event listing contains some events that are too small to warrant inclusions as SPEs under the SEC definition. It is important to note that a significant proton event will always be recognised whatever the definition, and it is only with 'low flux' proton events (with a peak flux of ~1 pfu) where there may be uncertainty as to the existence of an event. This type of 'low-flux'/poorly defined event is marked as class 3, so any uncertainty about the dataset can be eliminated by filtering out events of class 3 to leave only well defined events with good signatures.

One reason for including 'small' events as SPEs is due to the fact that the observed flux profile of an SPE is a function of observer location, hence a small SPE at earth may manifest as a larger SPE at a different location in the earths orbit. Also, it is possible that all SPEs, regardless of their peak flux, could have precursors of similar magnitude, hence failure to include 'small' SPEs may unnecessarily deplete the dataset.

In compiling the dataset events of class 4 or 5 were filtered out immediately. Events of class 2 and 3 were left in the dataset for potential removal at a later stage in the analysis. The exclusion of class 5 events is self explanatory. The exclusion of class 4 events was performed to try and restrict the dataset to the first instance of particle injections during a sequence of activity. SPEs that immediately follow an ongoing SPE could be due to interplanetary shock enhancement of existing solar protons rather than a separate injection. Such 'SPEs' are therefore unlikely to exhibit a correlation with other solar quantities (in particular x-rays), meaning that their inclusion in the dataset could corrupt any patterns in predictor variable behaviour. The total event list contained 284 SPEs occurring between 1965 and 1999.

To avoid comparing predictor variables from solar minimum with those from solar maximum (and thus detecting a solar min-max difference as opposed to an SPE / non-SPE difference) the study was restricted to SPEs occurring in solar active years as defined by Feynman et al. [*Feynman, Armstrong, Daogibner, and Silverman*, 1990b], these being the periods spanning from 2 years before to 4 years after the year of solar maximum of each solar cycle. Active year periods are given in Table 6-1. Table 6-2 lists the number of SPEs per category for solar minimum and solar maximum periods. It can be seen that the majority of SPEs occur during solar active years, so limiting the study to solar maximum periods does not significantly reduce the size of the SPE dataset.

Solar Cycle	Solar Maximum	Period of Active Years
20	1968.9	1966.9-1973.9
21	1979.9	1977.9-1984.9
22	1989.9	1987.9-1994.9
23	2001.2	1998.2-

Table 6-1 Times of solar maximum and periods of active years for solar cycles 20-23 [*Feynman, Armstrong, Daogibner, and Silverman*, 1990b].

Classification Type	**Number of Events** **Solar Max**	 **Solar Min**
1 – Suitably Well Defined	93	19
2 – Gradual Rise	20	1
3 – Poorly Defined	32	4
4 – Occurrence in Tail	4	1
5 – Unsuitable	33	2
Not Found (GOES satellites not operational)	66	9
Totals	248	36
Total	**284**	

Table 6-2 Breakdown of the total SPE list in terms of event 'class' as defined above.

The SPE listing is given in full in Appendix A.

6.2.3 Quiet Period Listing

The development of a binary classification technique required the collection of data pertaining to the two possible outcomes, thus it was necessary to compile a contrasting list of occurrence times corresponding to points in time at which no SPE occurred. These occurrences have been termed 'quiet periods' (QPs) and were defined as:

> *"Points in time at which the GOES >10MeV integrated proton flux was at a background level and had been so for at least 10 days"*

Quiet periods were generated by marking 'non-quiet' periods on a proton flux timeline spanning from 1974 to 1999 and taking points at random from the remaining periods. All quiet periods were from solar active years (as defined in Table 6-1) and were taken at least ten days from the end of an SPE. 340 quiet periods were generated in total. The figure is essentially arbitrary, but as more Quiet Periods occur in reality than SPEs it was deemed important to have a larger number of Quiet Periods than SPEs available for use in the study. A full listing of Quiet Periods is given in Appendix D.

6.3 Extraction of Precursor Data

With a list of times corresponding to the cases of SPEs and Quiet Periods predictor variables were collected to provide examples pertaining to each case.

6.3.1 GOES Solar X-Ray Data

The GOES satellites monitor solar x-rays in the 0.5-to-4 Å (XS channel) and the 1-to-8 Å (XL channel) wavelengths in units of Watts/m^2. GOES solar x-ray data exists as an off-the-shelf data product from 1986-1999, providing x-ray fluxes with a minimum resolution of 5-minutes. Archived GOES x-ray data was purchased to cover the period 1974-1986 and code was written to extract the data and process the 3-second raw values into 5-minute averages. X-ray flux data between 1974 to 1986 was found to be of poorer quality than the off-the-shelf data and contained several large data gaps and a higher level of noise.

X-ray data was extracted for time windows placed about the occurrence time of SPEs and QPs. C-code was written which took a list of times as an input, and then hunted through x-ray data, searching for the periods to extract. XS fluxes, XL fluxes (in Watts/m^2) and the ratio of the XS/XL channels were extracted for each SPE and QP in 5-minute resolution, and further C-code was written to process the data extractions into 24, 12, 6, 3 and 1 hour averages. The length of the extraction window typically spanned from –120 hours to +48 hours relative to the start time of each SPE and QP, although this could be easily altered to generate different extraction window lengths. Error checking algorithms were employed to ensure that each 5-minute average was based on at least 70% good data, and erroneous spikes in the data were removed (this was mainly used to increase the quality of the archived GOES x-ray flux data from 1974-1986).

Note that GOES x-ray data could only be extracted for SPEs occurring after January 1974 as prior to this GOES satellites were not operational.

6.3.2 Solar Radio Data

2800Mhz solar radio flux data was sourced from the NGDC web-site◆, consisting of daily averages from 1947 to 1999 in units of 10^{-22} $J.s^{-1}.m^{-2}.Hertz^{-1}$. Values for the radio flux up to 1991 were based on measurements from a ground based radio telescope near Ottawa. Since 1991 measurements have been based on observations from the radio telescope in Penticton, B.C., Canada. Adjusted radio flux values were used, which corrected for the variation in sun-earth distance.

Again, C-code was written to process a list of times and hunt through the radio data extracting radio fluxes for a time window centred on the occurrence time of each SPE and QP. Fluxes were extracted in daily average resolution for a period spanning from –81days to +81days relative to each occurrence.

6.4 Flare Associations

Findings by Reames and others which suggest that there are two groups of SPEs (as described in section 2.5) indicate that it may be useful to segregate SPEs according to their x-ray flare associations. Specifically, significant SPEs have been found to have a strong association with long duration X-ray flares, therefore it is possible that this type of event may have stronger precursors in x-rays than SPES associated with impulsive flares.

GOES x-ray flux plots were viewed to assess the temporal flare associations for each SPE. This was restricted to SPEs which occurred after January 1986 as prior to this GOES x-ray data were not available as an official data product and could not be viewed without significant processing effort.

◆ http://www.ngdc.noaa.gov/stp/SOLAR/FLUX/flux.html

If an SPE occurred within a few hours of an x-ray flare the flare was taken as being proton associated and was assigned as being either an impulsive flare or a long duration flare. Impulsive flares were defined as having a duration of minutes to 2 hours and took the form of an impulsive rise and decay in the GOES XL channel (1-8Å). LDX flares were defined as having a duration of >2 hours and were characterised by a steep increase in the XL flux followed by a gradual decay to background levels. Flare associations were made for 97 SPEs occurring in solar active years after January 1986, the breakdown of which is given in Table 6-3. A full listing is given in Appendix C.

Flare Association	Number of SPEs
Impulsive	30
Long Duration	60
No apparent flare association	7

Table 6-3 impulsive and long duration flare associations for SPEs occurring in solar active years after January 1986.

In addition to classifying x-ray flares as long duration or impulsive further x-ray flare associations were made according to the NGDC category of x-ray flare. The NGDC categorises flares according to the order of magnitude of the peak burst intensity (I) measured at the earth in the GOES 1 to 8 Ångstrom band (XL) as shown in Table 6-4.

X-Ray Flare Class	Peak burst intensity range (I) / W/m^2
B	I < 1.0E-06
C	1.0E-06 <= I < 1.0E-05
M	1.0E-05 <= I < 1.0E-04
X	I >= 1.0E-04

Table 6-4 X-ray flare categories as defined by the NGDC.

Flare listings of every monitored x-ray flare are freely available on-line and were downloaded for the years 1986 to 1999. By matching the occurrence times of SPEs to the occurrence times of x-ray flares on the NGDC lists, flare categories of proton associated flares were derived. For cases when several flares occurred around the start time of an SPE, GOES x-ray data were plotted and viewed by eye in order to identify the flare most likely to be proton associated. Flare associations were found for 98 SPEs occurring in solar active years after January1986 according to Table 6-5.

X-Ray Flare Class	Frequency
B	None
C	10
M	46
X	35
None	7

Table 6-5 X-Ray flare categories for 98 SPEs occurring in solar active years after January 1986.

A full listing is given in Appendix B.

6.5 Detrending of the Solar Cycle

Solar x-ray flux and solar radio flux vary significantly over the 11 year solar cycle, as shown below in Figure 6-1 and Figure 6-2, and although the study was limited to solar active years only, this 7 year period still exhibits a significant amount of variation due to the long term solar cycle.

The use of precursor data that contains the solar cycle trend may result in periods closer to solar minimum being compared to periods closer to solar maximum, hence there will be a difference in measured fluxes simply due to their relative positions in the solar cycle. This may mask the presence of any real SPE precursors, hence it was desirable to remove the solar cycle trend from the x-ray and radio flux data.

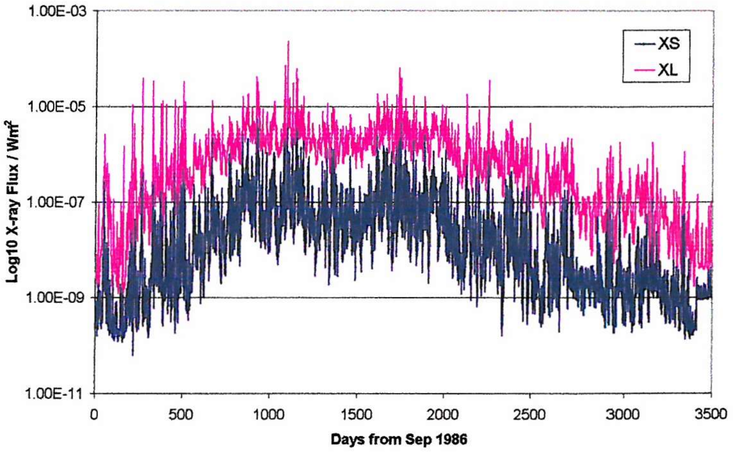

Figure 6-1 Variation in X-ray flux in the GOES XS and XL channels over solar cycle 22 (Sep86 – Apr-96).

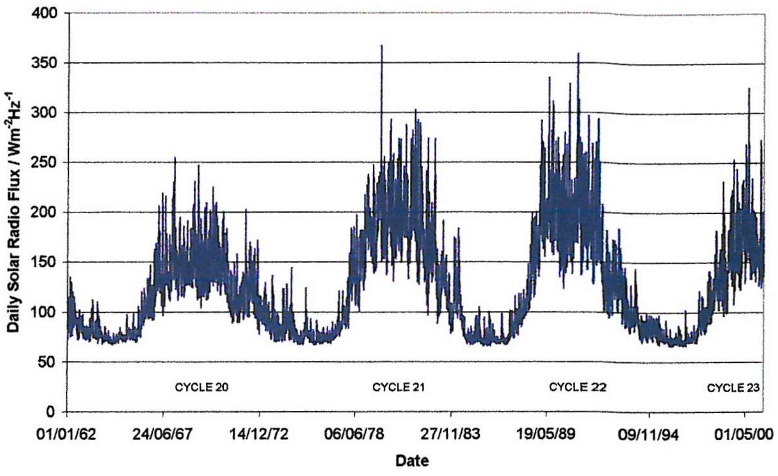

Figure 6-2 Variation in solar radio flux over cycles 20, 21, 22 and 23.

Limiting the periods of consideration to the 7-active years of the cycle reduces the visibility of this trend, but in order to eliminate it completely the trend over the 7 active years of each cycle was approximated by regression with a 2^{nd} order polynomial.

For each solar cycle, the trend function was subtracted from the absolute value of measured flux to produce a detrended value. Trend functions for each solar cycle are summarised below in Table 6-6.

Solar x-ray data from solar cycle 21 (1977.9-1984.9) was sourced from the archived GOES measurements and was found to consist of data from 6 different satellites, the first of which were of early design and relatively unreliable. Cross validation of measurements between the satellites revealed significant calibration differences between sensors (particularly during the SMS to GOES transition in 1976-1977), resulting in poor continuity of a uniform dataset. This, coupled with sparse and poor quality data from 1974-1979, rendered calculation of an accurate trend for solar cycle 21 impossible. Detrending of solar x-ray data was thus limited to measurements from post 1986, which only incorporated solar cycles 22 and 23.

Cycle 22 X-Rays

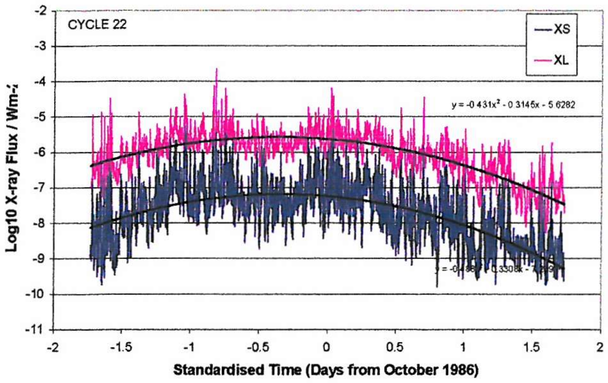

XS Trend

$y = -0.488x^2 - 0.3308x - 7.2391$

XL Trend

$y = -0.431x^2 - 0.3145x - 5.6282$

Cycle 23 X-Rays

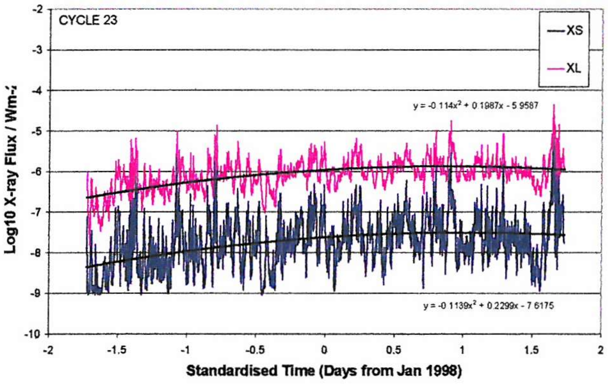

XS Trend

$y = -0.1139x^2 + 0.2299x - 7.6175$

XL Trend

$y = -0.114x^2 + 0.1987x - 5.9587$

Table 6-6 Trend functions for GOES solar x-rays over the active years of solar cycles 22 and 23. Cycle 23 is approaching maximum and is not yet complete.

Radio flux data was of consistent quality and solar cycle detrending could be carried out for all data post 1966 (i.e. the epoch at which the SPE listing began).

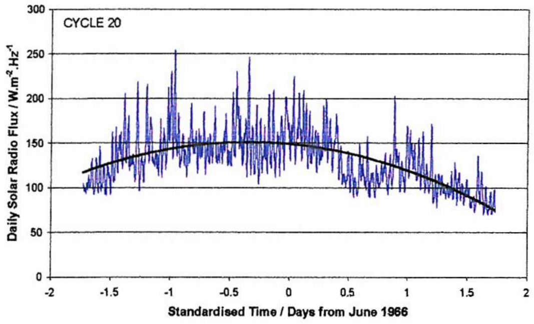

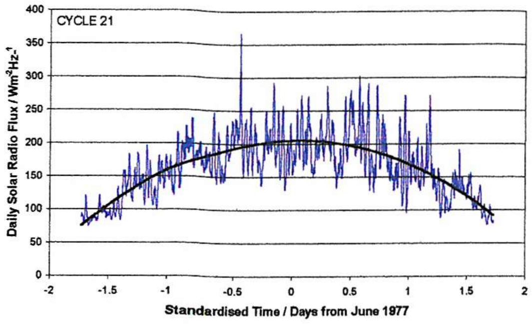

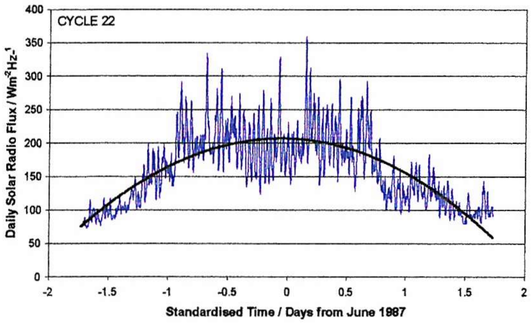

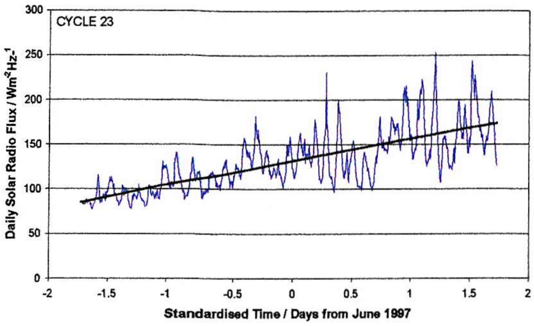

Cycle	Active Years	Equation
Cycle 20	1/6/66 – 31/12/73	y = -17.598x2 - 12.109x + 148.96
Cycle 21	1/6/77 – 31/12/84	y = -40.595x2 + 4.8356x + 205.41
Cycle 22	1/6/87 – 28/2/94	y = -46.714x2 - 4.8164x + 207.03
Cycle 23	1/6/97 – 31/12/99	y = -0.782x2 + 25.725x + 131.91

Table 6-7 Trend functions for 2800Mhz solar radio flux over the active years of solar cycles 20, 21 and 22 and 23.

6.6 Summary of Dataset

A list of 284 SPEs was derived from IMP and GOES solar proton data, spanning the period from 1965 to 1999. Any events that could not be identified due to missing data or due to their occurrence being associated with a recent SPE were removed from the dataset. To provide a contrasting case, a list of 340 Quiet Periods was generated relating to instances in time at which an SPE did not occur and the >10MeV proton flux had been at background for at least 10 days. Instances of SPEs and Quiet Periods were limited to solar active years to prevent the comparison of periods from solar maximum and solar minimum.

GOES x-ray fluxes in Watts/m^2 in the XS and XL channels and the 2800Mhz solar radio flux in units of 10^{-22} J.s^{-1}.m^{-2}.Hertz^{-1} were extracted at the times of SPEs and Quiet Periods. Solar radio flux was extracted for all events occurring since 1965, whereas GOES x-ray data could only be extracted for SPEs occurring after 1974.

For SPEs occurring after January 1986 flare associations were made, categorising the flare as impulsive or gradual. The NGDC flare category was also recorded and would enable the SPE listing to be segregated by flare association at a later stage.

7. DEVELOPMENT OF A CLASSIFICATION PREDICTION MODEL

This chapter describes the method used to develop an SPE forecast classification model based on the concepts detailed in Chapter 5. Classification models were optimised by altering input configurations, changing key parameters in the models and by selectively filtering the listing of SPEs. MLP neural networks were compared with RBF models and a linear classification technique.

7.1 Method

A default configuration of the classification model was arbitrarily adopted by defining an input window of length 72-hours and a lead time of 48-hours. Lead time was deliberately chosen to be significantly greater than that offered by current SPE forecasting models. The averaging period within the 72-hour input window was varied between 1 and 12 hours.

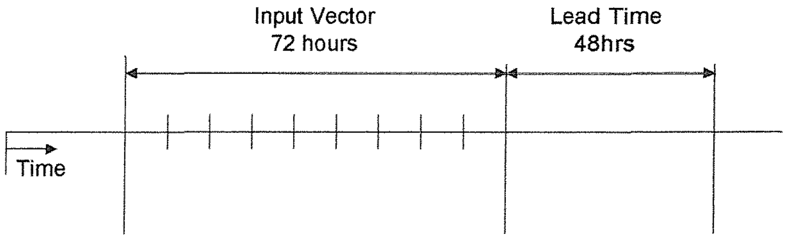

Figure 7-1 Default input configuration with a 72-hour input window and a 48-hour lead time.

The $\log_{10}$ of the ratio of the solar x-ray fluxes was used as the predictor variable input. Taking the ratio of the XS/XL fluxes combines the two values into one dimension, reducing the complexity compared to two separate input vectors, and the XS/XL ratio is also a measure of the spectral hardness of solar x-ray emissions. Garcia and Kiplinger have shown certain variations in spectral hardness during an x-ray flare to be indicative of an SPE [*Kiplinger*, 1995] [*Garcia and Kiplinger*, 1996].

MLP models were used to assess different input configurations due to the speed with which they could be constructed. The learning algorithm was halted when the training error stabilised, which typically occurred between 10,000 and 80,000 training epochs[◆] depending on the size of training set used. An example of training output from the network softeware is given in Appendix E. After determining an optimal input configuration MLP models were compared to RBF and Linear classification models.

7.1.1 Model Generation

Classification models were constructed using examples from the dataset described in the previous chapter. The total dataset was split into a training set and a test set. Neural models learned solutions on the training set and were then tested with the unseen test data. In cases where a Principal Components Analysis was used the derived statistics of the training set were used to transform the test set. A summary of the procedure is given in Table 7-1 and is shown schematically in Figure 7-2.

	Procedure for Model Generation
1	Divide dataset into a training set and a query set.
2 *	Perform PCA on training set
3	Train Model with training data
4 *	Apply previous PCA transformation to test set
5	Run Model with test data
6	Evaluate performance

Table 7-1 Summary of procedure for generating classification model. (*where applicable).

The number of examples of SPEs and Quiet Periods in the training set was always held equal in order to prevent biasing the training towards a particular outcome. For example, given 99 quiet periods and 1 SPE a 1% error can be achieved by simply classifying all examples as QPs, although clearly such an approach is not intelligent for SPE prediction. Using equal numbers of examples ensured that any learned solution would be due to some real difference between examples. Other options to this approach are addressed later.

[◆] 1 training epoch passes after each example in the training set has been presented once to the network.

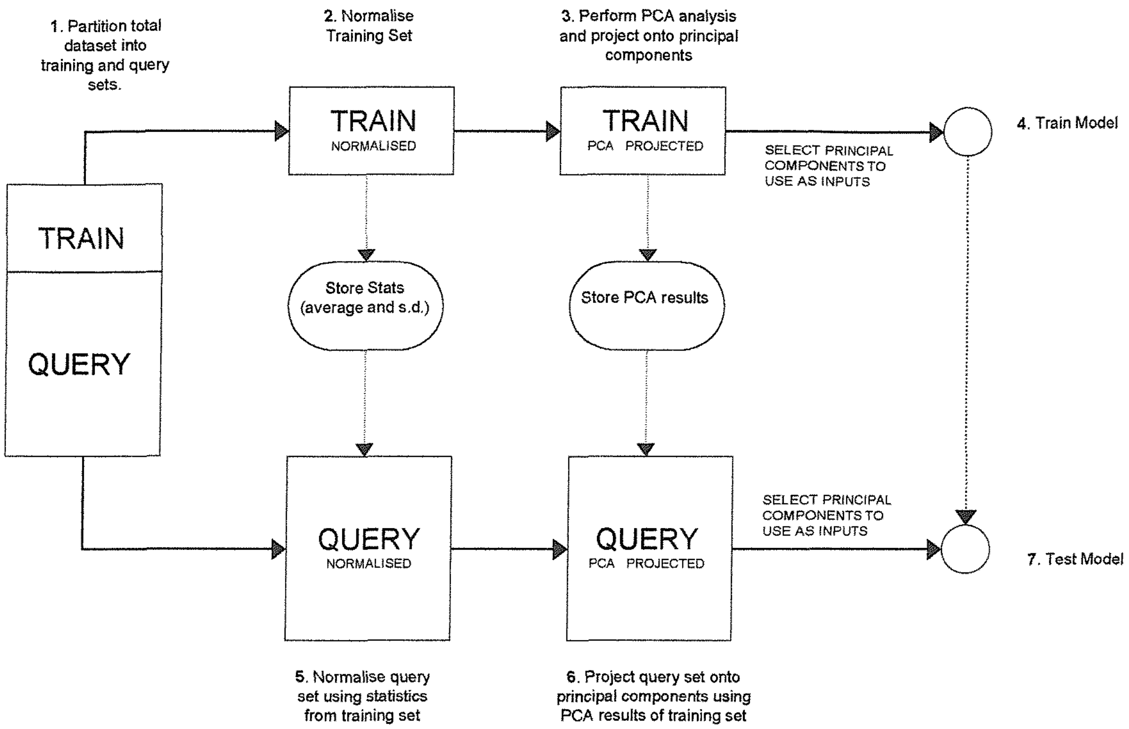

Figure 7-2 Schematic diagram showing the process for the construction of a classification model.

Software

MLP neural models were created using Neuframe v4.0. This is an off-the-shelf neural network package from 'NCS Manufacturing Intelligence' running on Windows and enables MLP networks to be created via a graphical user interface. Neuframe was run on a PC platform with a 500MHz Intel Celeron Processor.

RBF models were developed using software written by QinetiQ (formerly DERA). The software, denoted TSAR (Time Series Analysis Routine), was composed of a series of C-code modules and was command line driven in a UNIX environment.

Linear models were created in both MS Excel and MINITAB.

Input Scaling

Input scaling was performed within the respective neural network software packages, with the effect that each input is scaled to a value of between 0 and 1. To allow for the test dataset containing values slightly larger than the training dataset headroom of 10% was given in the training scale (effectively mapping the largest training value to 0.9 and the smallest to 0.1).

Training Targets

The supervised learning techniques employed by all models required target values for the two classes of SPE and QP. Input vectors corresponding to SPEs were given target values of 100 and input vectors corresponding to QPs were given target values of 0. During the model construction process target values were also scaled linearly to between 0 and 1, thus the choice of values to represent the two classes is arbitrary. Training targets of –1 and +1 for example will be scaled to the same training targets as 0 and 100. The Neuframe software also enabled 'text' training targets to be assigned to the input vectors instead of numerical values. Text targets allow discrete outputs to be specified as opposed to a continuous value.

Assessment of Performance

Models were assessed by analysing their response to each vector in the test set. Outputs of equal to or greater than 50 were interpreted as an SPE classification and outputs of less than 50 were interpreted as a QP classification. The effect of varying the decision threshold was addressed at a later stage.

The performance of a given model was measured by counting the number of examples in the test set that were classified correctly and then expressing this as a percentage of the test set. Note that there is only one degree of freedom in the model output as the model must generate one of only two possible outcomes. The number of SPE forecasts therefore also determines the number of QP forecasts.

Cross Validation

A cross validation technique was employed in order to ensure that a representative figure for performance was measured for each model. For a given configuration the training and test process was repeated 10 times, with a different combination of training and query data drawn from the same dataset. The performance for each model was measured as the average performance from the 10 variations. Cross validation measures the robustness of the model and takes into account the natural variation in performance from different training/query set combinations.

7.2 Results

7.2.1 Input Resolution

The time resolution of the input window for the default configuration was varied between 1-hour and 12-hours in order to determine the resolution giving optimum performance. The number of inputs to the model is dependent on the time resolution of the input vector, as shown in Table 7-2.

Time Average of the –120 to –48 hour x-ray ratio window.	Number of inputs to network
12 hour	6
6 hour	12
4 hour	18
1 hour	72

Table 7-2 The number of elements in the input vector for different time resolutions of the input window.

For each input resolution the size of the training set was varied between 30 and 240 examples (always keeping an equal ratio of SPEs and Quiet Periods) in order to find the optimum. This accounted for the fact that the number of free parameters (weights) within each network was also changing with the input resolution. Networks with a greater number of inputs were expected to require a larger training set in order to represent the larger number of possible input permutations. The combination of input resolution and training set size resulting in the highest performance were selected for further development.

Table 7-3 gives results for MLP neural models using input resolutions of 1, 4, 6 and 12 hours respectively. No data pre-processing was performed on the dataset other than taking logs of the values. Highlighted rows show the training set size giving the highest success rate for each input resolution. The network structure in Table 7-3 denotes the number of nodes in each layer of the model. The number of nodes in the first layer is dictated by the number of inputs to the model (which is dependent on the time resolution), and the number of nodes in the hidden layer is a default value chosen by the Neuframe software, based on the number of nodes in the input layer. Models were trained with discrete text targets of 'event' and 'quiet period' hence there are 2 nodes in the output layer.

12 HOUR AVERAGES (6 inputs, network structure of 6:3:1)

Training Set Size	EVENTS / %	QPS / %	OVERALL/ %
30	56.5 ± 7.5	55.3 ± 7.4	55.9 ± 1.7
60	54.5 ± 10.4	48.9 ± 7.0	51.7 ± 5.1
120	57.7 ± 12.0	45.4 ± 7.5	51.5 ± 5.1
180	54.0 ± 13.5	39.0 ± 7.8	46.5 ± 6.5
240	51.6 ± 19.6	30.3 ± 7.2	40.9 ± 12.3

6 HOUR AVERAGES (12 inputs, network structure 12:4:2)

Training Set Size	EVENTS / %	QPS / %	OVERALL/ %
30	53.7 ± 5.8	51.2 ± 6.7	52.5 ± 4.1
60	56.4 ± 9.4	51.4 ± 9.6	53.9 ± 3.9
120	54.0 ± 7.8	54.9 ± 6.1	54.5 ± 1.3
180	52.3 ± 10.4	53.8 ± 8.1	53.0 ± 3.8
240	46.6 ± 12.5	55.8 ± 7.0	51.2 ± 6.0

4 HOUR AVERAGES (18 inputs, network structure 18:5:2)

Training Set Size	EVENTS / %	QPS / %	OVERALL/ %
30	53.5 ± 11.4	49.8 ± 11.3	51.6 ± 4.3
60	57.0 ± 6.9	53.4 ± 7.7	55.2 ± 2.3
120	50.0 ± 3.7	57.6 ± 4.8	53.8 ± 2.3
180	51.0 ± 11.3	55.6 ± 6.0	53.3 ± 4.1
240	55.0 ± 9.6	53.8 ± 3.5	54.4 ± 3.9

1 HOUR AVERAGES (72 inputs, network structure 72:7:2)

Training Set Size	EVENTS / %	QPS / %	OVERALL/ %
30	48.0 ± 8.4	51.6 ± 10.1	49.8 ± 3.7
60	50.7 ± 8.5	55.2 ± 4.3	53.0 ± 5.0
120	53.0 ± 7.2	52.6 ± 9.2	52.8 ± 3.5
180	51.5 ± 6.6	54.0 ± 4.9	52.7 ± 3.6
240	58.4 ± 12.1	49.6 ± 6.6	54.0 ± 5.1

Table 7-3 Classification success for the default model configuration (72-hour input window, 48 hour lead time) for varying time resolutions and training set sizes.

Results from Table 7-3 are plotted below in Figure 7-3.

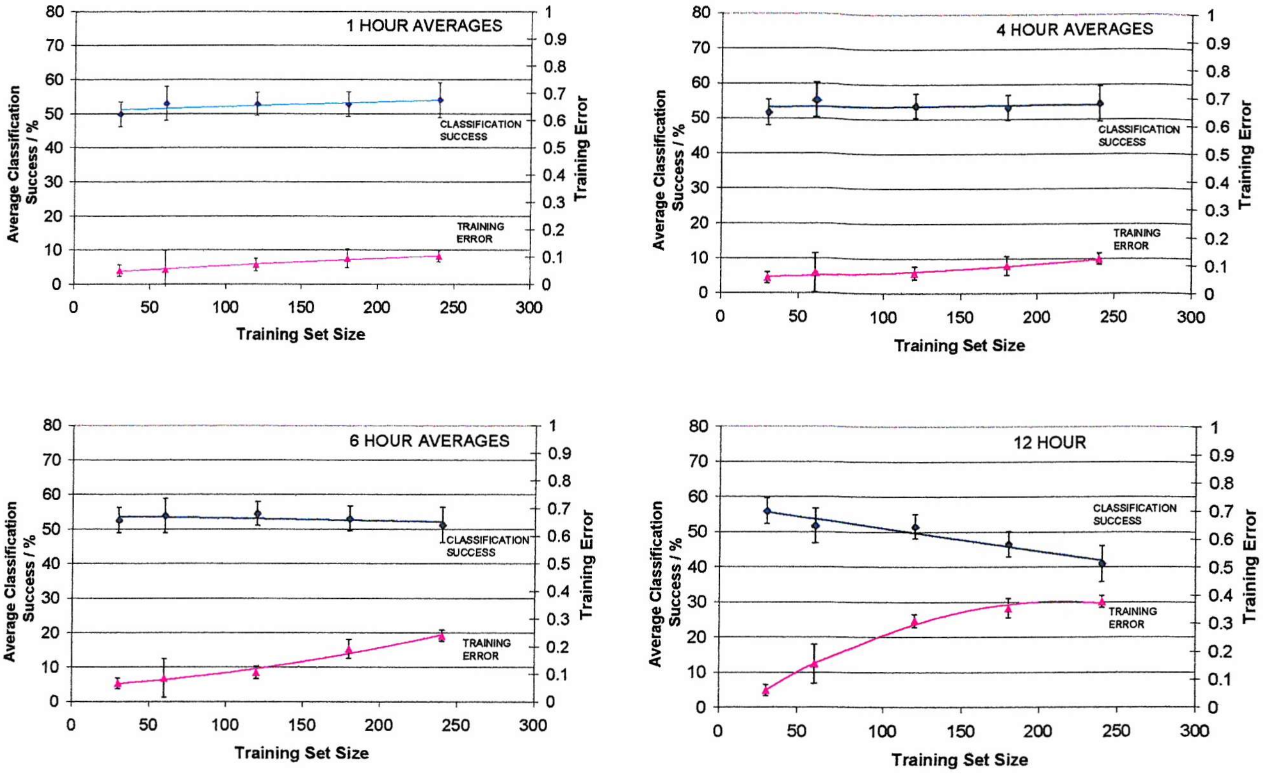

Figure 7-3 Plot of average classification success as a function of training set size for differing input resolutions. Training error as a function of training set size is also shown. Error bars denote 1 standard deviation.

Classification success for all models was measured to be between 50 and 55%. Given that if the model were to output 'event' and 'quiet period' at random one would expect a 50% classification success, the performance of the technique is not significantly better than random expectation. This would appear to indicate that the problem is ill-posed: i.e. there is little differentiation between the examples of quiet periods and the examples of SPEs in the dataset.

Performance was found to be independent of the size of the training set, implying that models trained as well on 30 examples as they did on much larger training sets of 120 and 240 examples. The only case in which this was not true was when using 12-hour resolution inputs, where larger training sets of greater than 60 examples were found to cause the model to generate a large proportion of 'undefined' outputs, resulting in low performance. These 'Undefined' outputs occurred as a result of using discrete class targets of ‘event’ and ‘quiet period’, and occurred when input vectors could not be satisfactorily mapped onto either of the two training classes.

As to why a 12-hour input resolution causes poor performance, the 12-hour input network has only 6 inputs and a small network structure with a low number of free parameters (weights). This means it is difficult for the model to approximate the variation observed within large training sets. The implication is that the training data contains no clear patterns or has contradictory examples which result in the model fitting a ‘neutral average’ that tends to map all input vectors to a similar output. This is supported by the high training error that is observed for 12-hour averages and large training sets.

Performance variation between models of different input resolution was not significant, being within 1 standard deviation of the cross-validation variation. However, the highest performance of the tested configurations was generated by 12-hour resolution inputs and a training set of 30 examples, hence this configuration was selected as the benchmark for further development. The advantage of using such a small training set is that a larger number of examples can be used to test the model, lending the test results a higher statistical significance.

7.2.2 Dataset Filtering

The SPE dataset was next filtered to study the effect of including only certain types of SPEs. In the first instance, class-3 (small) SPEs as defined in Chapter 6 were removed from the dataset in order to leave only significant SPEs of class-1 and class-2. In the second instance the SPE listing was filtered so as to include only SPEs associated with long duration x-ray flares. Thirdly, the x-ray dataset was detrended as described in section 6.5 to assess the effect of removing the long term solar cycle trend from the inputs. All dataset filtering was performed in conjunction with the optimal MLP model configuration determined in the previous section (7.2.1) – i.e. the model giving highest overall classification success. Results are given below in Table 7-4 and expressed graphically in Figure 7-4.

Model	Events / %	QPs / %	Overall / %
Benchmark 12 hour inputs.	56.5 ±7.5	55.3 ±7.4	55.9 ±1.7
Removal of class 3 events from dataset	54.5 ±11.0	53.3 ±6.1	53.9 ±6.2
Filtered list to include only events associated with long duration x-rays	58.6 ±11.9	56.7 ±7.6	57.6 ±5.9
Removal of solar cycle trend from XS/XL inputs.	58.2 ± 5.8	49.9 ± 5.5	54.0 ± 3.6

Table 7-4 Results from filtering SPE dataset to remove class-3 events and to leave only SPEs associated with long duration x-ray flares. Error denotes 1 standard deviation.

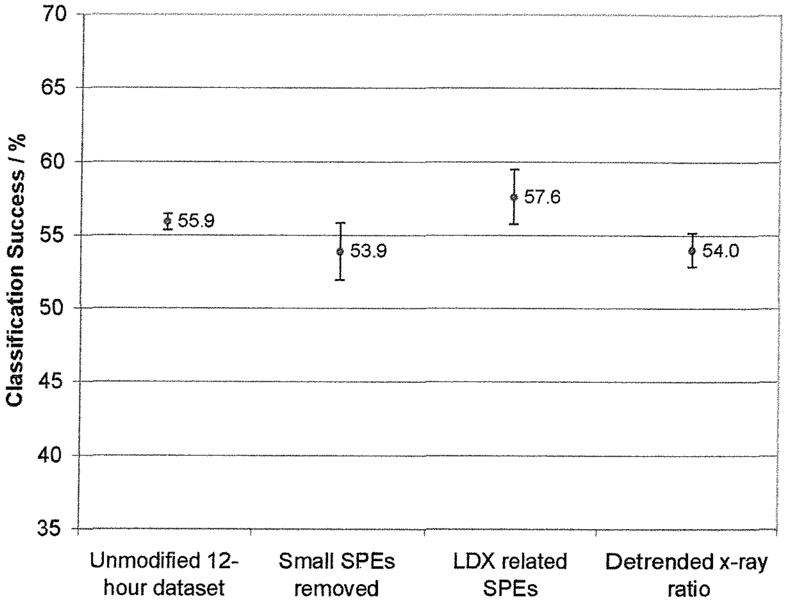

Figure 7-4 Graphical comparisons of different MLP model configurations shown in Table 7-4. Error bars denote the standard error.

The removal of small (Class-3) SPEs from the dataset had a slight but not significant detrimental effect on performance, indicating that the inclusion of small SPEs in the dataset was not a cause of general poor performance. By showing that the classification ability of the technique is not improved by limiting the dataset to well defined SPEs it is also implied that well defined SPEs are not preceded by well defined precursors (at least during the input window being considered).

Filtering the dataset so as to include only SPEs associated with long duration x-ray flares marginally increased classification success by 1.7% over the benchmark model. Whilst not significant, the result suggests that SPEs associated with LDX flares are more distinguishable from quiet periods than a group of SPEs with mixed flare associations. This implies that SPEs associated with LDX flares may have better defined precursors in the XS/XL ratio during the input window. The reason why this filtering is successful probably stems from the fact that the dataset has been filtered with respect to the behaviour of the *predictor* variable (i.e. it is guaranteed in the filtered dataset that all examples of SPEs coincide with long duration x-ray flares). Although the input window does not include the x-ray flare itself, it is composed of a value derived from x-ray flux measurements, and it is known that a flare does occur, hence the model may be detecting precursors to the flare. By comparison, when filtering out small SPEs from the dataset no consideration is given to the predictor variable and it is unknown whether or not a significant x-ray flare occurs at the time of the SPE.

The detrending of a time series prior to analysis is a common practice, designed to isolate local fluctuations from long term trends. Removing the long-term solar cycle trend from the XS/XL ratio was found to produce a 54.0% classification success. This is 1.9% lower than the benchmark model. The result suggests that the presence of the solar cycle trend in the inputs may have been slightly beneficial. One explanation for this is that SPEs have a general association with the solar cycle. Knowing that the XS/XL ratio is higher during solar maximum than solar minimum, the model could classify 'high' value input windows as SPEs and 'low' value input windows as quiet periods, and via the general association between SPEs and solar maximum the strategy would probably identify more than 50% of cases correctly. The fact that detrending the inputs has still resulted in a comparable performance to trended inputs indicates that the technique was not simply using a solar cycle association in order to function, and suggests that real SPE/non-SPE differences may be present.

7.2.3 Multiple Model Configuration

As an experiment, a novel approach was used to combine the outputs from different models into one prediction with the aim of increasing performance. This configuration has been termed a multiple-model. The technique created 10 different models from 10 different random choices of training data and then queried each model simultaneously with the same set of test data. This caused 10 responses for each example in the test set, and a 'majority decision' was made as to the overall outcome. Figure 7-5 below shows a schematic diagram of the multiple model format.

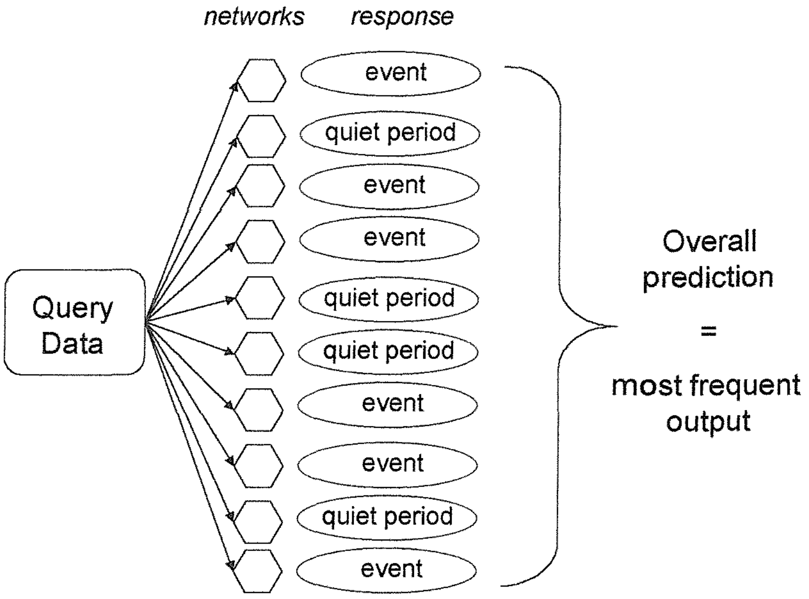

Figure 7-5 Schematic diagram showing configuration of the 'multiple model'.

Multiple model configurations were created for the filtered SPE listings described above, and their performance compared against that of the average 'single' MLP model. Note that the multiple model required 10 models to be trained for 1 test set, hence the build time was considerable and required a significant amount of data preparation. For this reason a maximum of two multiple models were created for each SPE listing. Whilst this does not rigorously assess the variability in performance due to dataset variation, it does serve to indicate the approximate variance in performance that can be expected.

Results are shown below in Figure 7-6. The average performance of each of the individual networks composing the multiple model is also given to show the improvement offered by the multiple model configuration.

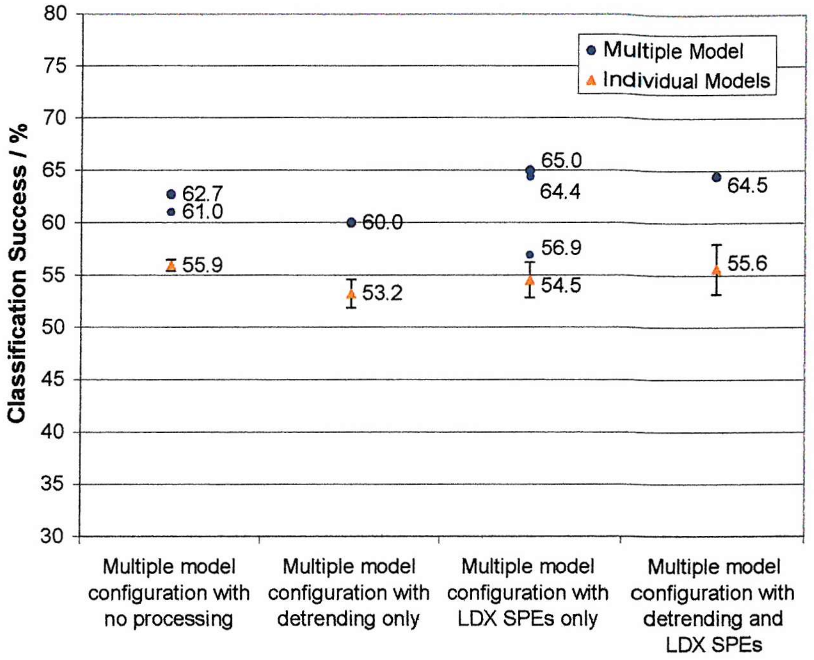

Figure 7-6 Multiple model performance in conjunction with dataset filtering and detrending. Points show the difference between multiple model performance and the average performance of the constituent MLP networks.

Multiple models of the standard configuration (12-hour inputs, no pre-processing or filtering) were found to generate a 61.0% and 62.7% classification success respectively. Given that the average individual network for the same configuration has only a 55.9% average success rate this indicates that a significant performance benefit is gained from simultaneously querying several models with one input vector and taking a majority response as the overall classification. The multiple model was assessed with a test set composed of 50 SPEs and 100 QPs, and a chi squared analysis showed that a 61% classification success was significant to a 99% confidence level. This demonstrates to a high degree of uncertainty that the model output is not random. An example of the Chi-Squared calculation is given in Appendix F.

Using an SPE dataset with only LDX flare associated SPEs has much the same effect with the multiple model as with single network models in that performance is increased from a non-filtered dataset by the order of 2%. The use of a detrended dataset also has a comparable effect as in the single models, reducing performance by around 2% from the benchmark. It can be seen that the effects of filtering and detrending the dataset are cumulative to the benefit of a multiple model configuration.

The effect of the multiple model configuration is to increase the classification ability of the technique from the 55% success level to the 62% success level. The results from duplicate multiple models indicate that the high success rates are reproducible and are not simply due to 'lucky' training-query set combinations. In all cases the performance of the multiple model using a majority vote is superior to the average performance of its constituent individual networks.

Whilst the presence of a solar cycle trend in the input data appears to offer a performance benefit it may be causing the technique to operate differently to how it was intended. Instead of using a difference between the cases of SPE and quiet periods to make classifications, models may be using an association between solar maximum and SPEs, causing high value input windows to be mapped to 'SPE' classes. The general association between SPEs and solar maximum could then result in a greater than 50% success rate, even though the technique is not actually detecting specific SPE 'precursors'. To ensure that the technique could only use SPE/non-SPE differences to assign classes further development was limited to detrended inputs, in which the long term solar cycle had been removed.

A decision was also made to use datasets that contained all SPEs, rather than only LDX flare related SPEs. Whilst a dataset containing only LDX flare associated SPEs results in higher performance, it also restricts the scope of the model by implying that it can only predict SPEs associated with long duration x-ray flares. This goes against the core requirements that were set during the conceptual design of the model, where it was stated that the model should not be dependent on the occurrence of a specific solar phenomenon. It is possible, for example, to have significant SPEs that are associated with large impulsive flares [*Garcia*, 1994a], but these would be excluded from the study if only LDX flare related events were considered. By encompassing all occurrences of SPEs in the dataset the model is 'universal' in its treatment of SPEs and is not confined to predicting SPEs of one origin.

A multiple model achieved a performance of 60% using detrended data and a non-filtered dataset (i.e. a dataset containing all SPEs, regardless of association), therefore a significant classification success is still achievable even with these 'constraints' in place.

7.2.4 Principal Components Analysis

The Principal Components Analysis (PCA) technique was described in detail in Chapter 5. The technique allows the dataset to be collapsed into fewer dimensions without significant loss of information, meaning that the number of inputs to the model can be reduced (reducing complexity) whilst at the same time increasing the signal to noise ratio.

The use of a PCA required the test set to be projected using coefficients derived from the training set. For this reason one would intuitively expect the success of the technique to be greater when the training set is more representative of the test set, and this is more likely to be the case when the training set is large.

This supposition is contrary to the findings in section 7.2.1 in which no PCA was used where smaller training sets were found to be as good as larger training sets. To determine the effect of training set size a further training set size optimisation was performed in conjunction with a principal components analysis.

A PCA was performed on a model using 12-hour resolution inputs, detrended values of the XS/XL ratio and a dataset containing all SPEs. Models were created using training sets of 30, 60, 90 and 120 examples. After determining the optimum training set size for 12-hour resolution inputs, the same training set size was used for an equivalent model using 3-hour resolution inputs. Results are given below in Figure 7-7.

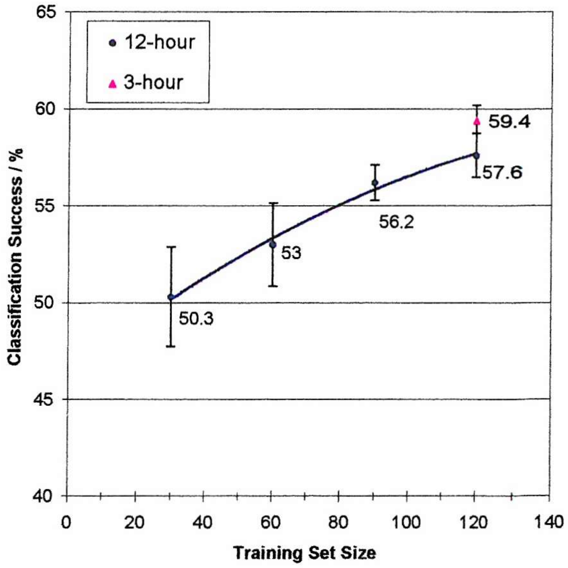

Figure 7-7 Classification success as a function of training set size for MLP models using principal components analysis projection of the dataset. Error bars denote the standard error.

It is clear that performance increases with the size of training set. Maximum classification success is achieved with the largest training set of 120 examples. 3-hour resolution inputs were found to be superior to 12-hour resolution inputs for the optimum training set size.

The use of a principal components analysis has resulted in a significant performance increase over the same model with no PCA. Figure 7-4 shows that a MLP network using detrended 12-hour inputs and no PCA has a 54% average classification success, whereas it can be seen from Figure 7-7 that the same model with a PCA has a 57.6% classification success. This is increased further to 59.4% when 3-hour resolution inputs are used with the PCA. It can also be noted that these figures are for individual networks, therefore a further performance increase can be expected if a multiple model configuration were to be used.

Further scope exists for increasing the success of the technique by optimising the number of principal components that are used as inputs to the model. It may only be necessary to use a small number of principal components to express the useful variation within the dataset. The 3-hour resolution model had a total of 24 principal components, and the number of these used as inputs was varied in stages between 3 and 24. Results are given below in Figure 7-8.

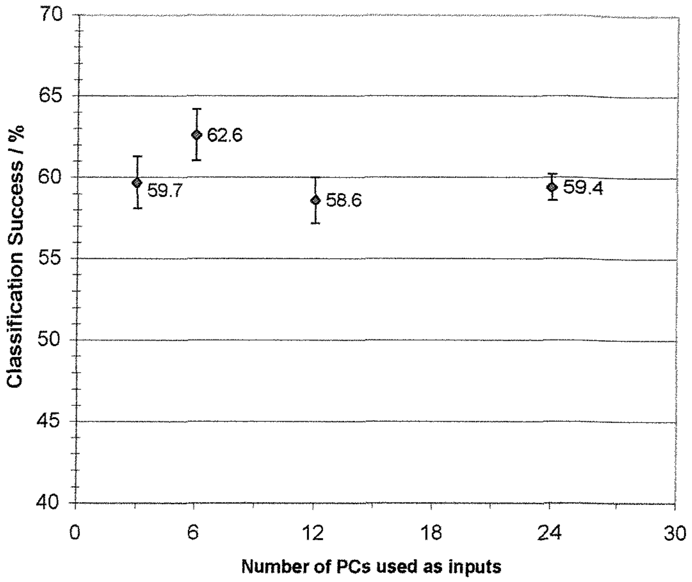

Figure 7-8 Classification success as a function of number of principal component inputs for MLP network using 3-hour resolution inputs. Error bars denote standard error.

Optimal classification success was generated when using the first 6 principal components of the 24 as inputs. With 3 inputs it would seem that not enough information is given to the model, whereas with more than 6 principal components it would appear that the extra input channels do not contain useful information.

7.2.5 Lead Time

The lead time of the technique was varied between 6 and 64 hours by moving the input window with respect to the time of the SPE or quiet period. This was performed for the optimal PCA model described above, i.e. an MLP model using the first 6 principal components as inputs from a 3-hour resolution input window. Results are plotted in Figure 7-9.

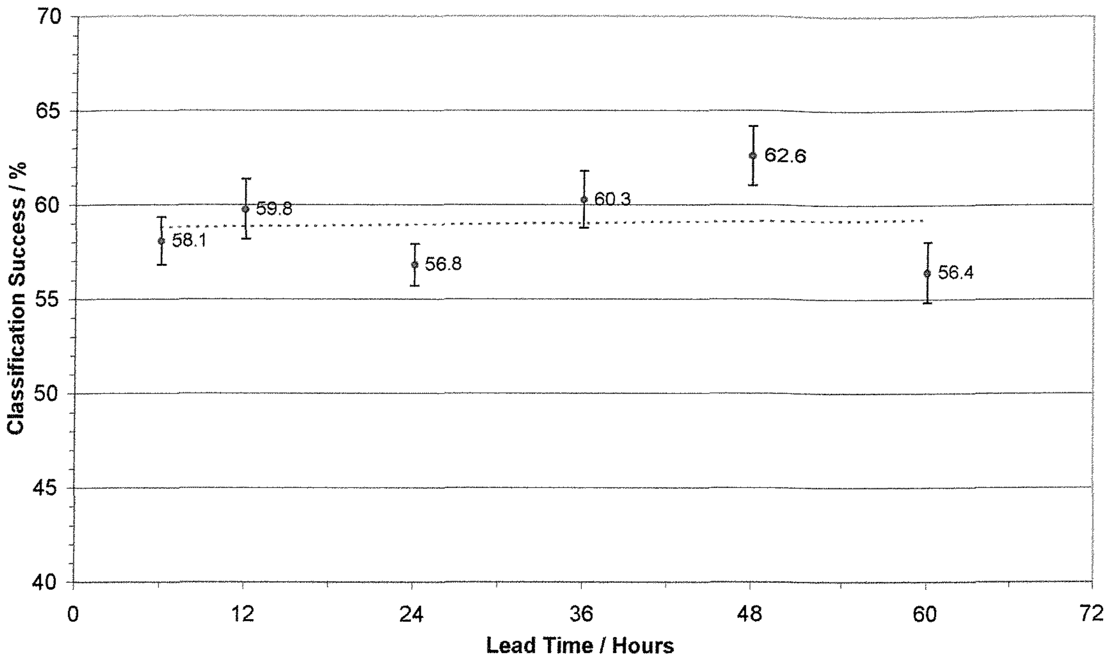

Figure 7-9 Variation in classification success as a function of lead time for the optimal PCA MLP model. Error bars denote standard error.

Classification success remains roughly constant regardless of lead-time, and surprisingly does not increase when the input window is moved significantly closer to the time of the outcome. Maximum classification success was found to occur with a 48-hour lead time, but would seem to drop away significantly if the lead time is increased further.

Figure 7-9 would appear to indicate that there is no benefit in moving the input window closer to the time of the SPE or quiet period, although this is likely to change if the inputs were brought forward so that they included the x-ray flare itself that is associated with the SPE. This would of course defeat the object of a long lead-time model.

7.2.6 Standardisation of the Input Window

The use of a second order polynomial to approximate the solar cycle variation results in the solar x-ray flux being zero meaned over the 7 year active period of solar maximum. This means that over a short period of time, e.g. 72-hours, the solar x-ray flux may have a non-zero mean dependent on the current level of solar x-ray activity with respect to the current solar cycle trend average.

As an alternative method of detrending, models were developed based on individually standardised input windows. Each input vector was standardised by subtracting the average and dividing by the standard deviation, resulting in all training and test examples being zero-meaned.

The performance of the two detrending methods was compared using MLP models in conjunction with a PCA and 3-hour resolution inputs of the XS/XL ratio. Results are given below in Table 7-5.

Model	Overall / %
Solar Cycle Detrended	62.6 ±5.0
Row Normalised	51.8 ±4.6

Table 7-5 Comparison between solar cycle detrending and row normalising.

It can be seen that independent row normalising produces a classification success that is significantly worse than solar cycle detrending, and is not in fact significantly better than random (i.e. 50%). The result shows that the classification ability of the technique must stem from the variation of the XS/XL with respect to the solar cycle average.

7.2.7 Radial Basis Function Models

Radial Basis Function neural network models were constructed to act as a comparison to the MLP networks. Models were based on the optimal configuration identified previously for the MLP models, i.e. a 48-hour lead time, detrended XS/XL inputs and a principal components analysis.

Before comparing MLP and RBF models, the parameters of the RBF model were optimised. The training set was fixed at 120 examples (60 SPEs and 60 Quiet Periods), and the number of centres was varied incrementally between 10 and 120. A thin plate splines function was used as the centres function. This was performed for 12-hour and 3-hour resolution inputs. Results are plotted below in Figure 7-10.

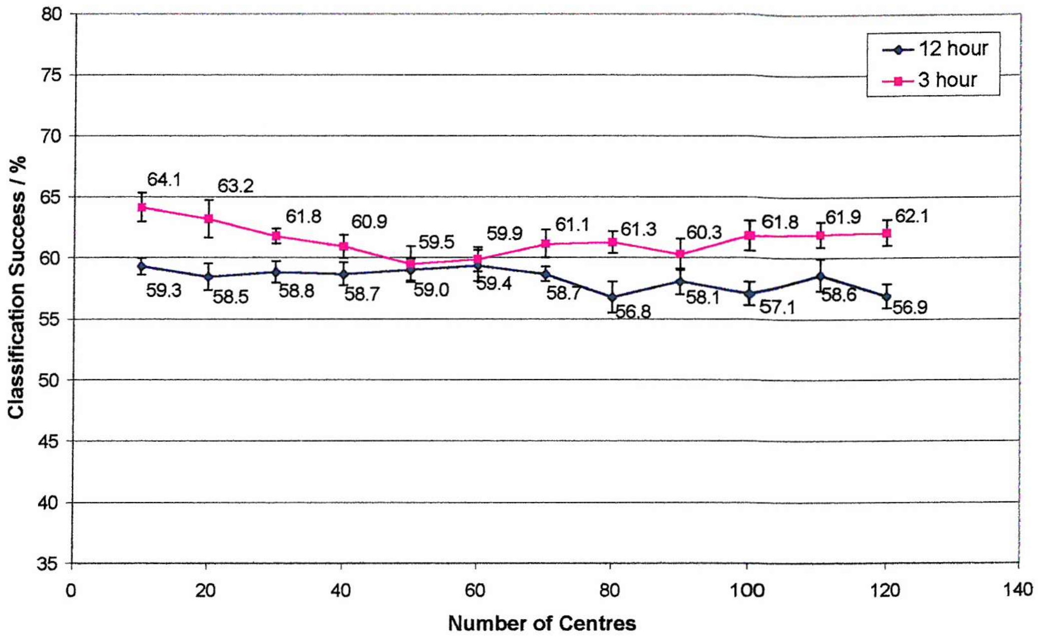

Figure 7-10 Variation in classification success with number of centres for RBF model using a principal components analysis and detrended XS/XL inputs in 3 and 12-hour resolution input windows. Error bars denote standard error.

As for the MLP models, 3-hour resolution inputs were found to give consistently superior performance to 12-hour resolution inputs.

Optimal classification success was given by using just 10 centres in the RBF model. No benefit is gained by a more complex fitting surface from a larger number of centres. This is probably due to the use of the k-means clustering algorithm for centres placement. This places the centres functions in their most efficient locations in the network subspace, meaning that additional centres functions offer little benefit. Although increasing the number of centres will have resulted in a more complex fitting surface, the results from Figure 7-10 show that the generalisation is worse, suggesting that additional centres functions are only fitting noise.

Centres Function Optimisation

The type of centres function was optimised using those available within the TSAR software. RBF models using 3-hour resolution inputs and 10 centres were created for 8 different centres functions. Results are plotted below in Figure 7-11.

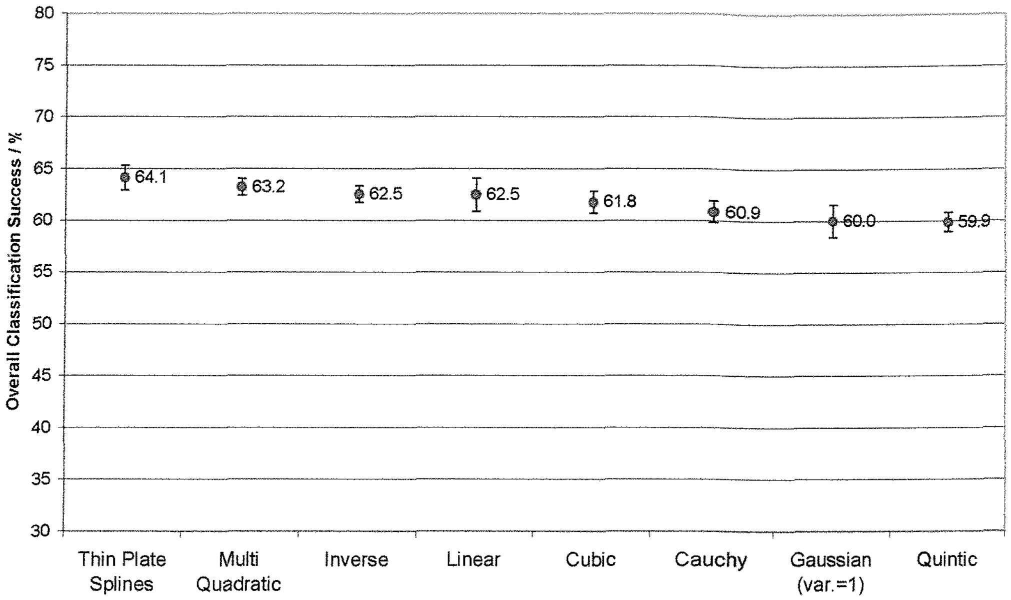

Figure 7-11 Classification success for different centres functions using an RBF model, arranged in order of effectiveness. Error bars denote standard error.

No major difference was found between centre function types, although the best performance was generated by the thin plate splines function: a function of the form:

$$\varphi(r) = r^2 \log r$$

(in which 'r' is the Euclidean distance between a point in space and the point at which the centre function is placed). Using the thin plate splines as the centres function the RBF model was found to generate a 64.1% classification success.

7.2.8 MLP, RBF and Linear Model Comparison

The MLP and RBF models developed above were compared with an equivalent linear model. All used 3-hour resolution inputs, detrended data, and a principal components analysis. Figure 7-12 shows the optimisation of each model for the number of principal component inputs.

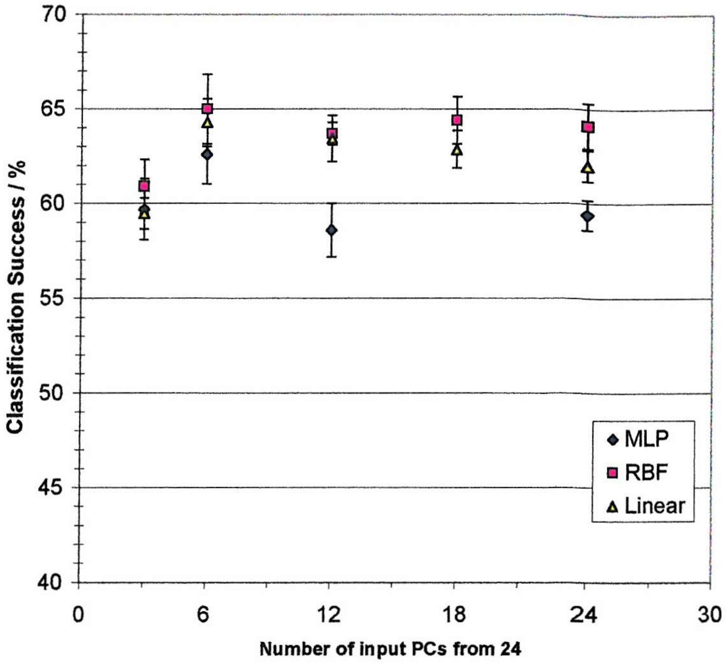

Figure 7-12 Classification success as a function of number of principal component inputs for MLP, RBF and Linear models using 3-hour resolution inputs. Error bars denote standard error.

For all model types optimal performance is given with 6 principal components as inputs. The optimal model of each type is compared below in Figure 7-13.

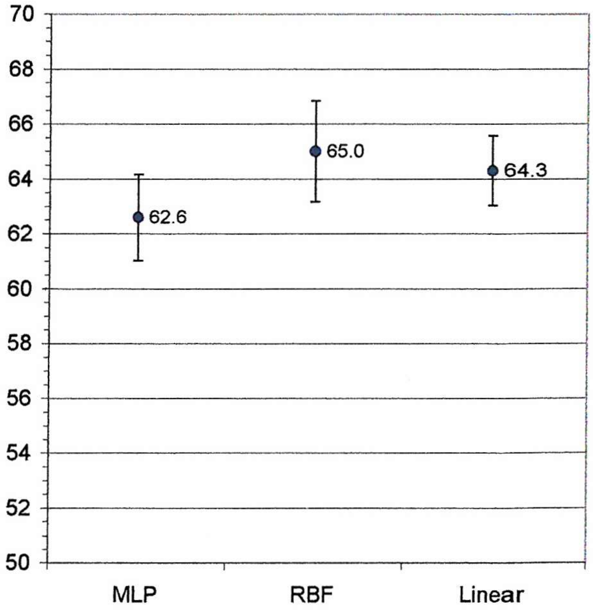

Figure 7-13 Comparison of optimal models (which use the first 6 PCs as inputs) for MLP, RBF and linear model types. Error bars denote standard error.

The value of the standard error means that the difference between model types is not significant, therefore the choice of classification technique is not a performance driver. RBF models generated the highest performance, closely followed by linear models. MLP models were found to produce the lowest performance.

One reason for the success of the RBF technique could be the relative ease with which its structure was optimised, in that it was possible to perform an optimisation for the number of centres. This is relatively easy in the case of the RBF model because the network structure contains only one hidden layer. The number of nodes in this layer is the number of centres.

No equivalent optimisation was performed in the case of the MLP model, as this would have meant varying the number of hidden nodes in the network. The problem with this is that both the number of hidden layers and the number of nodes per layer are variable, resulting in a potentially limitless number of network structures.

The fact that a linear model performs comparably to the non-linear techniques implies that the classification technique is not non-linear, or at least can be well approximated by a linear model. Whilst one can postulate that solar mechanisms are probably non-linear, the problem posed here is far abstracted from a physical process, and is simply trying to group input vectors as belonging to one of two classes, therefore it is perhaps not surprising that the linear technique performs well.

The fact that all model types perform to a similar level indicates that the success of the technique is primarily due to the content of the dataset. This means that SPE precursors may well be present in the XS/XL ratio.

7.3 Summary

This chapter has presented the first development stage of the classification technique, which was restricted to the use of the XS/XL ratio as inputs. Goals were to identify the effects of filtering and detrending the dataset, and to compare the performance from different model types.

Filtering the dataset so as to include only SPEs related to long duration x-ray flares was found to marginally improve classification success of initial models. This filtering was not adopted for further development though as it was felt that it compromised the validity of the technique by imposing the pre-requisite of an x-ray flare.

Detrended data was identified for use in the development of the models to ensure that the solar cycle variation in the inputs was not being used as a means to distinguish between the cases of SPE and quiet period. A comparison between solar cycle detrended inputs and locally normalised inputs revealed the latter to have very low performance. This indicated that the variation of the input with respect to the value of the local solar cycle average was important in determining the classification of an input vector.

A principal components analysis was found to be effective with large training sets, and gave superior performance with 3-hour resolution inputs. MLP, RBF and Linear models were all found to generate superior performance when using the first 6 principal components as inputs. The optimal configuration, irrespective of model type, consisted of 3-hour averages, the detrended log of the XS/XL ratio as inputs and a principal components analysis.

The difference between model types was not found to be significant, although the best performance was generated by the RBF network, followed by the linear and MLP models. These generated average classification success rates of 65.0%, 64.3% and 62.5% respectively.

A multiple model configuration was found to significantly increase the classification ability of the approach for all input configurations and SPE filtered listings. The performance of a multiple model was always found to be greater than the average performance of its constituent networks, suggesting that the PCA optimised RBF model may generate even higher performance than 65% when acting in a multiple model configuration.

The apparent success at predicting SPEs 48 hours in advance, even though limited, is surprising considering that no discrete x-ray flare is used as an input. This suggests the existence of SPE precursors on a timescale of days prior to SPE occurrence – a finding that has not been documented elsewhere.

The following chapter examines the behaviour of the solar x-ray ratio, the XS and XL fluxes and the solar radio flux, and compares the performance of models using these different variables as inputs. A quantitative analysis of the behaviour of the XS/XL ratio may help to explain the behaviour of the model and identify specific differences between the cases of SPEs and quiet periods that prove the existence of new SPE precursors.

8. PREDICTOR VARIABLE ANALYSIS

A classification technique has been shown to have limited success for SPE and Quiet Period prediction when using $\log_{10}$ of the XS/XL ratio as an input from between –120 and –48 hours relative to the outcome. This chapter examines the behaviour of solar x-rays for the SPE and Quiet Period case by plotting the XS and XL fluxes as a function of time, and identifies points at which the distributions are significantly different using the F-statistic Similar analyses are performed for the 2800Mhz solar radio flux with the aim of identifying any behaviour that can be associated with SPE occurrence. MLP classification models have been created using the different predictor variables as inputs in order to determine their relative performance.

8.1 Method

Detrended x-ray fluxes from the XS and XL channels were extracted in 3-hour resolution for a time period spanning –720 hours (27 days) to + 48 hours relative to each quiet period and SPE in the dataset. At each 3-hour interval the average value of the flux was calculated for the SPE examples and quiet period examples, allowing an average flux-time plot to be made for each of the two cases. X-ray plots were based on 97 SPEs and 192 quiet periods occurring between 1986 and 1999 inclusive. (The full listing of SPEs is given in Appendix C.)

A similar extraction was performed using daily resolution x-ray fluxes, for a period spanning –81 days to +81days relative to each event. 81 days is approximately equal to 3 solar rotations, and allowed flux-time plots to be generated over a longer time period for the SPEs and quiet periods.

Daily resolution 2800Mhz solar radio flux was also extracted for ±81 days relative to each SPE and quiet period occurrence. The existence of solar radio data from 1940 onwards allowed data extractions for the full list of SPEs, such that plots were based on 253 SPEs occurring between 1965 and 1999 and 340 quiet periods chosen at random between 1977 and 1999.

At each point in time the F-statistic was calculated to measure the statistical significance of any difference between the SPE and quiet period averages. The F-statistic is defined as:

$$F = \frac{MS_b}{MS_w} \qquad \textbf{8-1}$$

where MS_b is the mean square difference between the distributions, and MS_w is the mean square difference within each distribution. Given k distributions, with each distribution containing n_j observations, the sum of squares between (SS_b) and the sum of squares within (SS_w) are defined as:

$$SS_b = \sum_{j=1}^{k} n_j (\overline{X}_j - \overline{X})^2 \qquad \textbf{8-2}$$

$$SS_w = \sum_{j=1}^{k} \sum_{i=1}^{n_j} (\overline{X}_{ij} - \overline{X}_j)^2 \qquad \textbf{8-3}$$

The mean squares between and the mean squares within each distribution are thus calculated as:

$$MS_b = SS_b /(k-1) \qquad 8\text{-}4$$

$$MS_w = SS_w /(N-k) \qquad 8\text{-}5$$

where N is the total number of observations across all distributions.

The value of the F-statistic measured the statistical significance of the difference between the SPE and quiet period means and allowed points in time at which the two distributions were most different to be identified.

8.2 Results

8.2.1 Behaviour of Solar X-rays

Figure 8-1 displays flux time plots for the detrended XS and XL x-ray fluxes and the XS/XL ratio for –720 hours to +48 hours relative to each SPE and quiet period. Figure 8-2 plots the F-statistic as a function of time for the XS, XL and XS/XL distributions, and shows the statistical significance of the difference between the SPE and quiet period averages.

It can be seen from Figure 8-1 that the x-ray fluxes and ratio for the SPE case have a significant peak at close to time zero due to the x-ray flares that the SPEs are associated with. The average x-ray ratio and fluxes for the SPE case appear to rise from approximately 240 hours prior to event occurrence, whereas no similar increase is present in the quiet case. A study of the F-statistics in Figure 8-2 shows the two distributions to be significantly different in the means to a 99% confidence level from approximately 140 hours prior to occurrence. A rise in solar x-ray flux over several days prior to SPE occurrence has not been previously documented and would appear to be a new result.

It can be noted that the models developed in the previous chapter used inputs from -120 hours to -48hours, and this does appear to be a prime period during which the SPE and quiet period distributions are significantly different.

A comparison between the XS, XL and XS/XL ratio channels in Figure 8-1 shows a more consistent separation of the distributions in the XS and XL fluxes than in the ratio. In particular, quiet periods in the ratio show a greater level of variation than in either the XS or XL channels, and this appears to be why the F-statistic is lower for the x-ray ratio at some points close to the event occurrence time (e.g. at approximately –90 and –50 hours).

To some extent Figure 8-1 also indicates a small peak in the SPE x-ray flux at around 700 hours prior to occurrence. This is approximately 27 days (the solar rotation period) prior to occurrence, and may indicate that the activity at time zero is related to a recurring active region. A study of the F-statistic shows that the degree of separation between the two distributions at around –700hours is significant, and is strongest in the x-ray ratio. It would generally appear that the x-ray ratio is more volatile than the XS or XL fluxes.

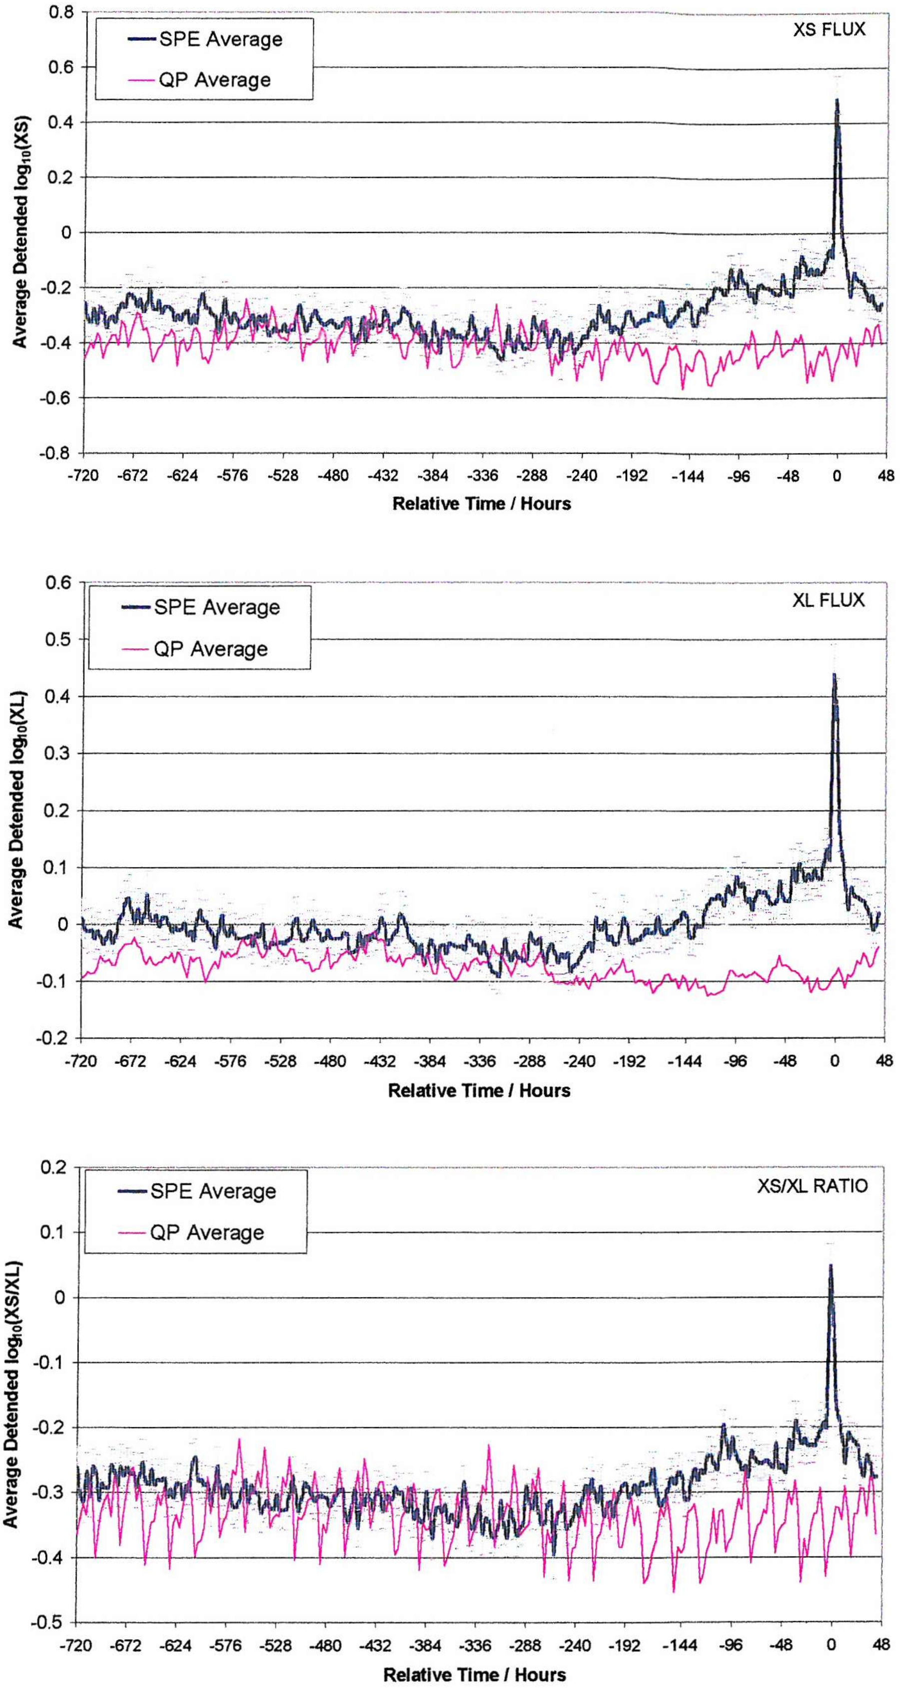

Figure 8-1 Average value of detrended XS, XL and XS/XL x-ray ratio prior to SPEs and quiet periods. Error bars denote standard error, and for clarity are only displayed for the SPE case.

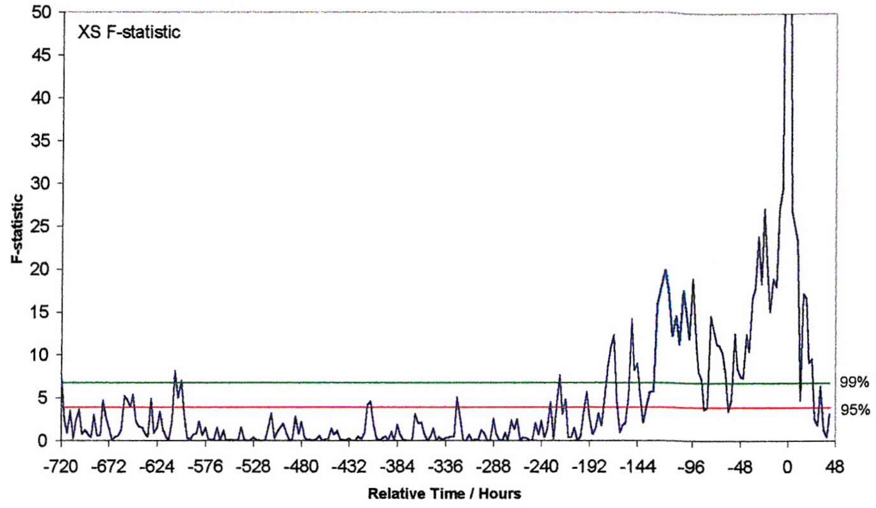

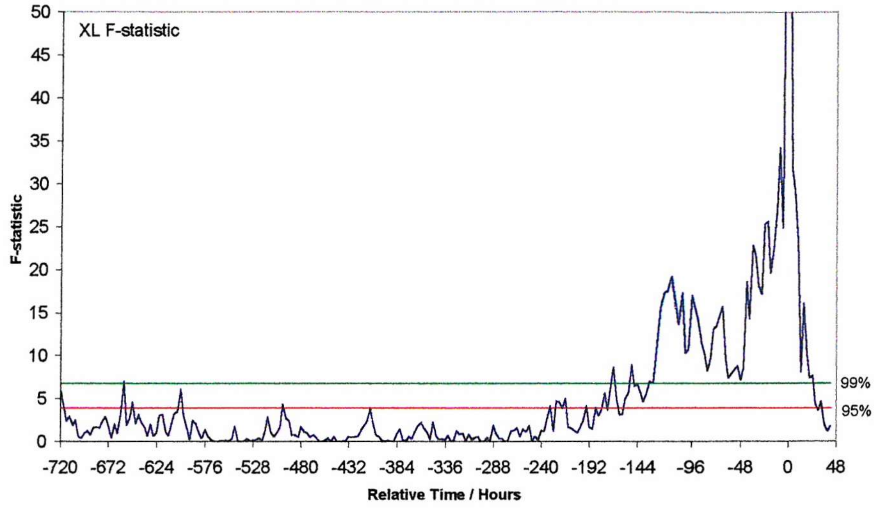

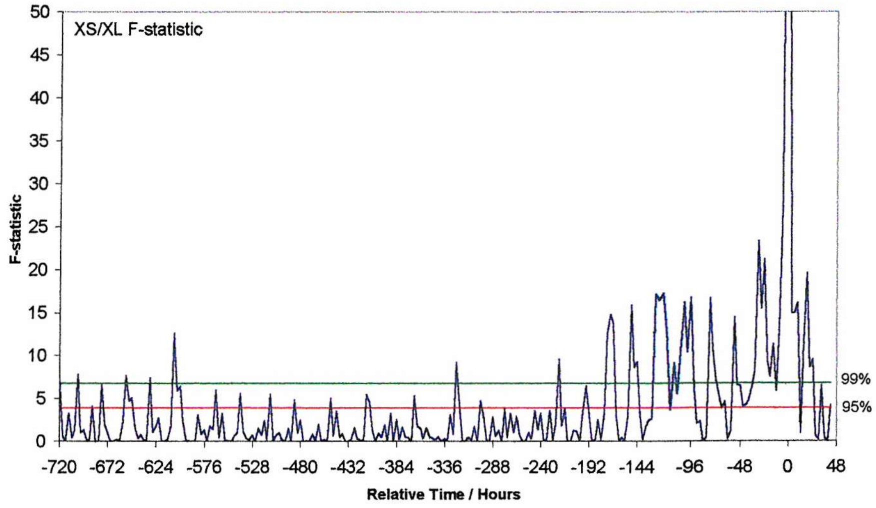

Figure 8-2 F-statistic plotted as a function of time for the XS, XL and XS/XL ratio showing the significance of any difference between the SPE and quiet period average.

The variation in x-ray intensity as a function of central meridian distance is presented by Donnelly et al. and shown in Figure 8-3. It shows that for a Central Meridian Distance (CMD) angle of less than 90 degrees the measured x-ray flux is relatively independent of the x-ray source location [*Donnelly and Puga*, 1990]. 100% of x-ray emissions will reach the observer right up until the source has rotated past the solar limb, at which point attenuation occurs rapidly.

The GOES x-ray detectors measure x-ray flux from the entire solar disk, but 98% of this flux is known to be contributed from active regions [*Wagner*, 1988]. As 98% of the solar x-ray flux is composed from active regions, and as the solar x-ray flux is not attenuated until sources reach the solar limbs, the GOES x-ray flux is predominantly a measure of active region emission over time - a fact that does not seem to have been stated in any other study. This means that the x-ray flux behaviour in Figure 8-1 is predominantly due to changes in active region emissions rather than the changing CMD of the x-ray source. The behaviour of the solar x-ray flux prior to SPEs, i.e. a general increase in flux over time, could therefore be explained by the notion of a developing active region that reaches an activity peak at the time of an SPE. The GOES measurements cannot of course distinguish between contributions from active regions that are present on the solar disk at the same time.

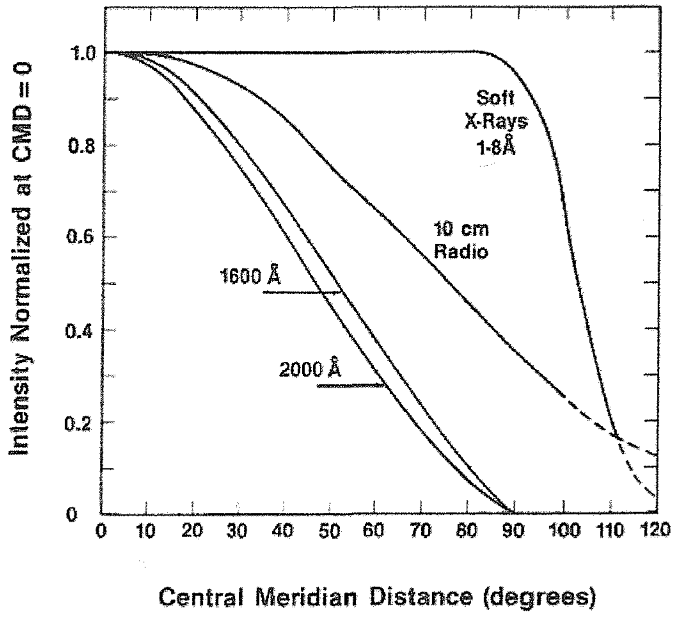

Figure 8-3 Dependence of transmitted solar emissions on the central meridian distance for different wavelengths as presented by Donnelly et al. [*Donnelly and Puga*, 1990]

Figure 8-4 shows the variation in the SPE and quiet period averages over a time span of ±3 solar rotations (81 days) relative to the time of occurrence. The relatively sudden increase in x-ray flux for the SPE case at around –10 days could be due to the reappearance of an active region on the eastern limb (at a CMD angle of ~90°) which would rapidly switch from 0 to near 100% transmission based on the dependence shown in Figure 8-3.

It can be seen that there is an apparent peak in the SPE case for all 3 plots, centred at around –27 days, again supporting the theory that activity at time zero is related to a rotating active region. There is no significant difference between distributions at greater than one solar rotation prior to occurrence however, indicating that any recurrence in x-ray activity is limited to the most recent solar rotation. The separation between the SPE and quiet period averages is similar in all of the x-ray channels.

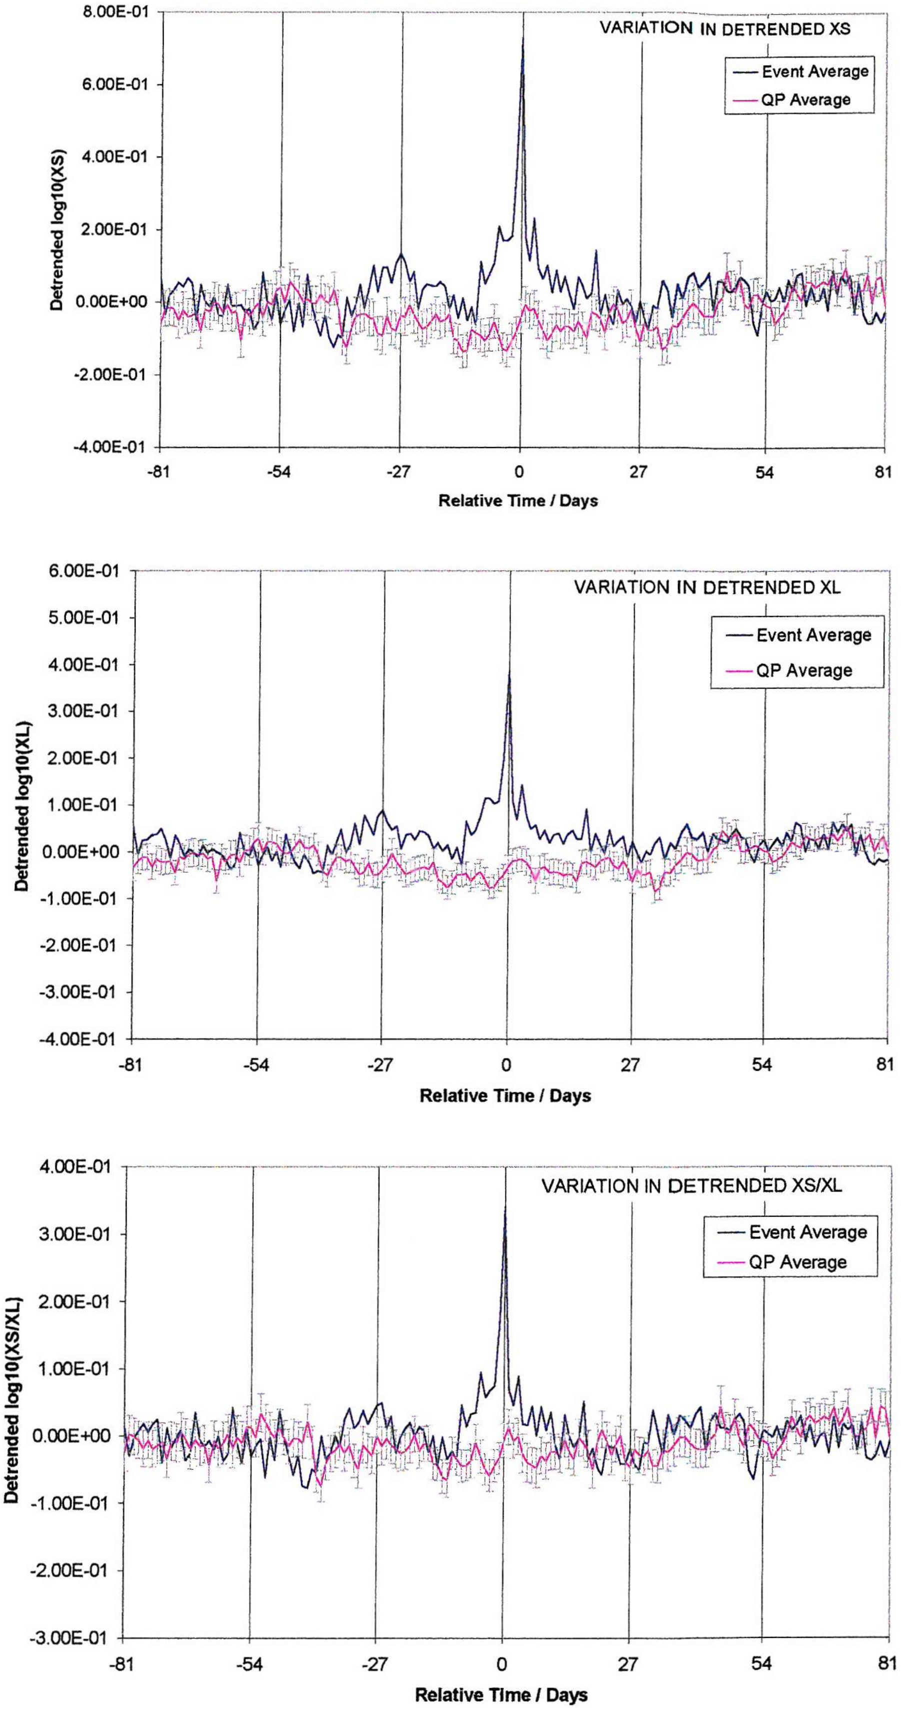

Figure 8-4 Variation in average XS, XL and the XS/XL ratio for SPEs and quiet periods over the period –81 to + 81 days relative to each event. Error bars denote standard error and for clarity are only plotted for the quiet case.

8.2.2 The Behaviour of Solar Radio Flux

Figure 8-5 plots the SPE and quiet period averages for the 2800Mhz solar radio flux between ±81 days relative to event occurrence time. F-statistics for the distribution are given in Figure 8-6.

The average radio flux associated with SPEs exhibits a strong 27-day periodicity, passing through a maximum at time zero (the time of the SPE) and at 27-day intervals either side of time zero. The periodic fluctuation is clearly visible for 2 solar rotations prior to the event to 1 solar rotation post event. The F-statistic (Figure 8-6) near to the –52, -27, 0 and +27 day points shows the distributions to be different in the mean to a greater than 99% certainty. The SPE radio peak at time zero is larger than the neighbouring peaks. SPE occurrence has not previously been associated with peaks in a 27-day periodicity in the solar radio flux, and this would appear to be another new result.

The radio flux associated with quiet periods also exhibits an approximate 27-day periodicity. Figure 8-7 displays the same data after normalising each example in the dataset, and rules out the possibility that the peaks are simply due to a handful of large examples weighting the 'SPE case' average. After normalising, distinct peaks are still seen at ± 27 days from the time of an SPE, and the flux associated with non-SPE examples is near a minimum at time zero. The high values of the F-statistic in Figure 8-6 indicate a significant separation of the distributions at certain points in time, and suggest that the radio flux may be a good predictor variable for the classification model.

The behaviour of the SPE radio flux could be explained by the concept of a long lived active region which develops and decays over several solar rotations, producing an SPE when it is at a maximum activity level. By contrast, quiet periods may be associated with times at which there are no significant active regions on the earth facing solar disk, corresponding to a minimum in the averaged radio flux. Note that after +27 days the SPE and quiet period averages are not significantly different, which could be due to the decay of SPE-associated active regions, and/or the development of active regions at previously 'quiet' heliolongitudes. The idea of a developing active region being associated with an SPE is supported by findings from Mursula and Zieger, who show that solar stream structures (from the magnetic fields of active regions) have a lifetime of around 4 solar rotations [*Mursula and Zieger*, 1996]: this lifespan is directly consistent with the SPE-case of Figure 8-5 that exhibits four peaks between –54 and +27 days, and certainly suggests that active regions are the cause of the periodicity.

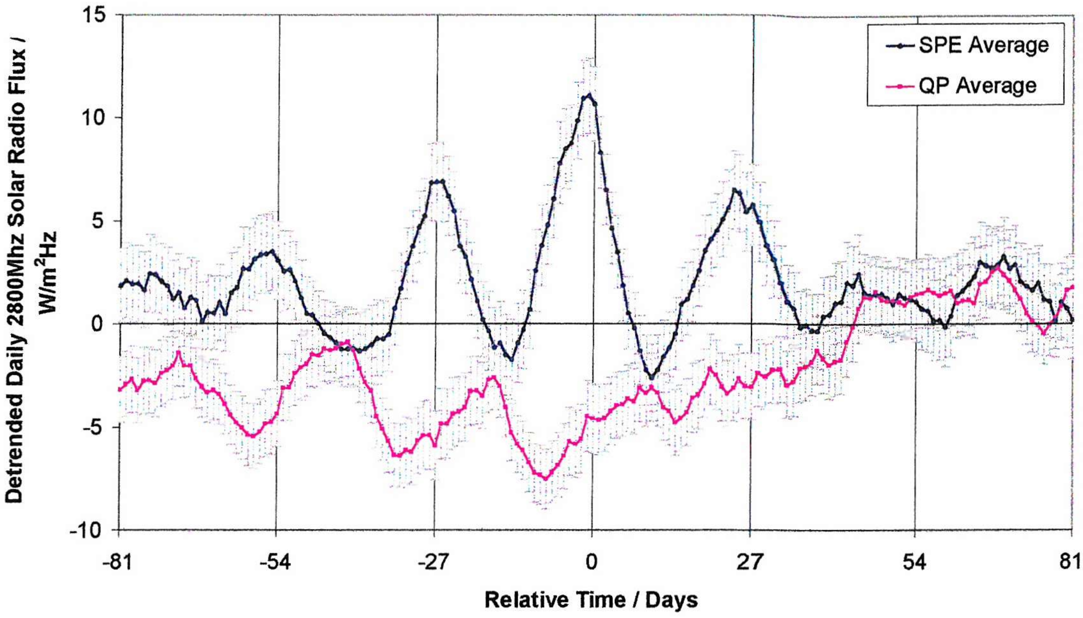

Figure 8-5 Variation in average detrended solar radio flux for SPEs and QPs for –81 days to +81 days relative to each event. Error bars denote standard error.

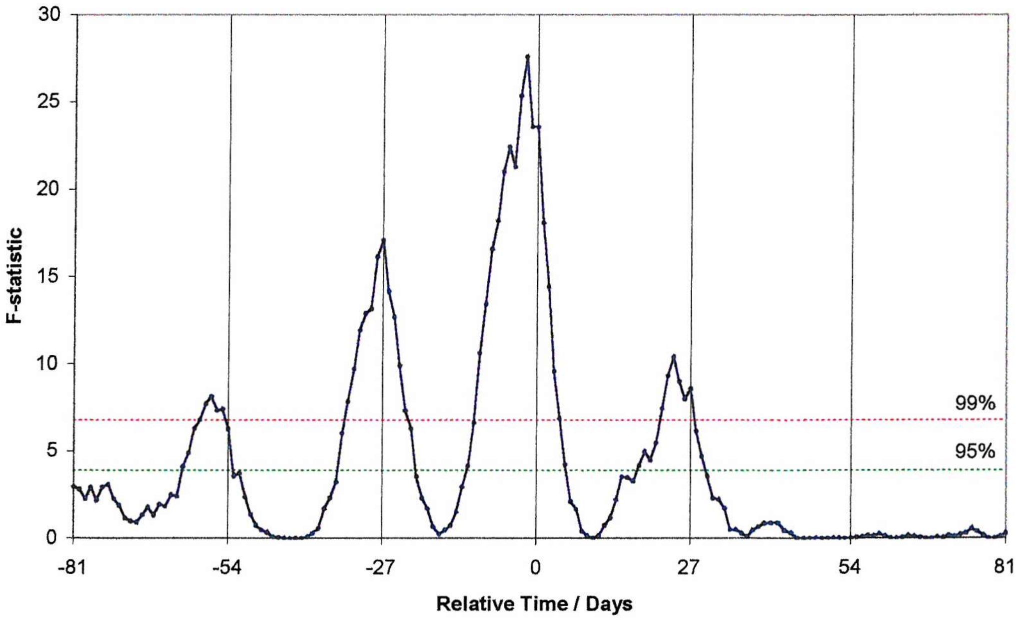

Figure 8-6 F-statistic plotted as a function of time for solar radio data, showing the statistical significance of the difference between the SPE and QP average.

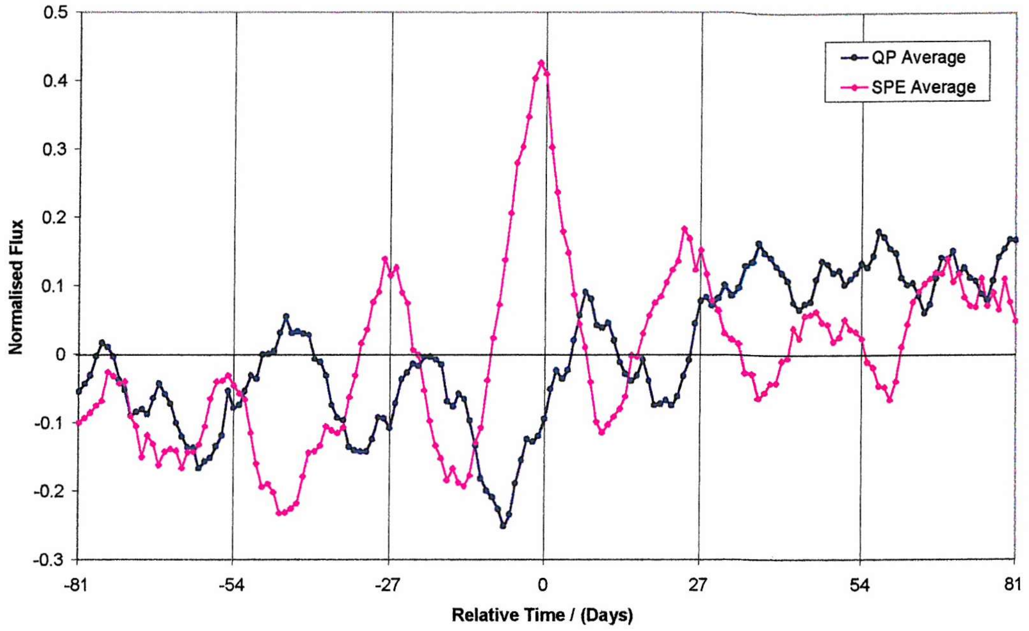

Figure 8-7 Variation in average detrended solar radio flux for SPEs and QPs after normalising each example in the dataset.

The variability in the solar radio flux over the 27-day solar rotation can be partly explained by the optical thickness properties of the solar atmosphere shown in Figure 8-3. A constant power radio source will be attenuated significantly as soon as it drifts beyond 30 degrees of the central meridian distance, and will be attenuated by a factor of 0.5 by the time it reaches the solar limb. Clearly, a 27-day periodicity can be explained by a constant power source in conjunction with a varying CMD angle.

However, the radio flux plots in Figure 8-5 show the average solar radio flux to be at a maximum at the time of an SPE, when it is highly unlikely that the active regions producing the SPEs will be at a central CMD location. (Shea and Smart [*Shea and Smart*, 1994] show that most SPEs occur west of 60° heliolongitude, indicating radio attenuation factors of at least 0.65 from Figure 8-3).

If radio flux maxima at time zero are assumed to emanate from predominantly limb-situated active regions it means that the observed variation in solar radio flux must be due to variations in the emitted flux as well as the varying CMD of the source location. The fact that a radio flux peak is seen even though the emitted flux is probably attenuated by a factor of ~0.65 indicates that at the time of an SPE the emitted flux is high enough to negate the increased attenuation from the solar atmosphere. Studying Figure 8-5, the central peak is placed approximately 2-days to the left of time zero. This could be explained by an active region that reaches the solar central angle ~2 days prior to time zero (giving maximum transmitted flux) and then produces maximum emitted flux (associated with an SPE) at time zero after it has passed the central CMD.

With the ability to measure radio flux spatially from a single active region alone a correction could be applied for the observer location angle which may yield a better correlation between radio flux peaks and SPE occurrence. The current radio flux data contains a variation due to changing CMD angle and a variation due to the source emission flux changing, but these are currently inseparable because measurements are not spatially resolved. Given that a radio flux peak is seen despite a probable attenuation of ~0.65 the emitted flux must be very high at that time, therefore removal of the observer location dependency may reveal a better defined peak in the radio flux that is highly correlated with the occurrence of an SPE.

The existence of radio flux peaks at 27-day intervals prior to and after the occurrence of an SPE provides new evidence for a strong association between SPEs and the development of an active region. The finding is in direct agreement with Chakravorti et al, who studied the relation between active region characteristics and SPE occurrences based on 171 proton events observed between 1966 and 1984 [*Chakravorti, Das, Sen, and Dasgupta*, 1991]. The Chakravorti study plotted the daily solar radio flux as a function of time for ±7days relative to the occurrence of an SPE and found the radio flux from active regions to pass through a maximum on the day of an SPE in 70% of cases. The study did not however consider the behaviour of the radio flux over longer periods, and the results of Figure 8-5 are the first time that SPE occurrence has been shown to coincide with peaks in a 27-day periodicity in the solar radio flux.

Figure 8-5 does of course plot the 'average' case for the SPE examples. In order to better quantify the correlation between SPES and radio flux peaks, the time between a local radio flux peak and the occurrence time of an SPE was found to the nearest day for 149 SPEs occurring after 1986. Data was binned and the histogram in Figure 8-8 shows the percentage of SPEs in the dataset as a function of the time from any local radio flux maximum. Figure 8-9 plots the same distribution as a cumulative frequency graph.

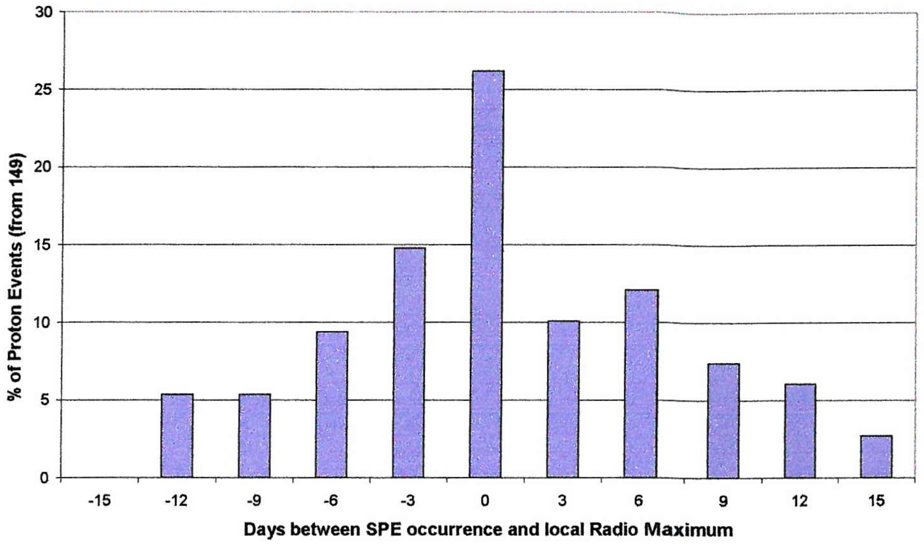

Figure 8-8 Histogram showing percentage of SPEs that occur within 'X' days of a local radio flux maximum.

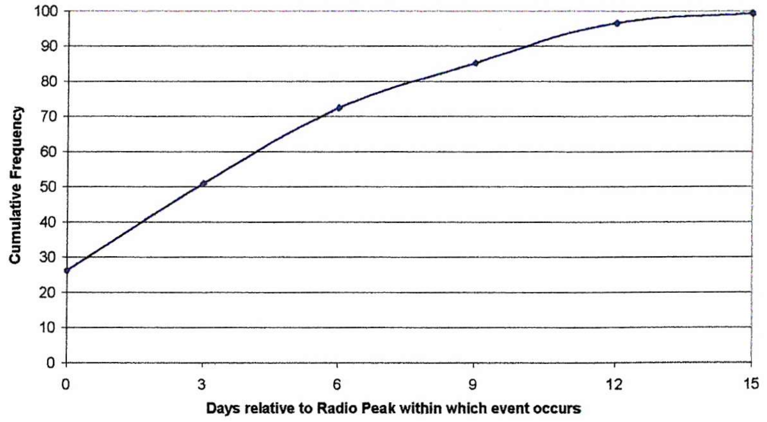

Figure 8-9 cumulative frequency plot showing percentage of SPEs that occur within 'X' days of a local radio flux peak.

It can be seen that over 25% of SPEs occur within ±1.5 days of a local radio flux peak, and 51% of SPEs occur within ±4.5 days of a local radio flux peak (taking into account the bin widths). The correlation is not as high as that reported by Chakravorti et al. (70%), but still shows a clear correlation (the difference in results may be due to the fact that Chakravorti used spatially resolved measurements from individual active regions, although the study is not specific in its exact source of data). If the radio flux peaks belong to a 27-day periodic fluctuation, then this means that over half of the SPEs occur within a time frame of 4.5/27*100 = 17% of the period. Theoretically, knowing when the last radio flux peak occurred it would be possible to calculate when the next peak will occur, and assign an SPE risk of occurrence at this time. This assumes however that a 27-day periodicity continually exists in the radio flux, and evidence for this is inconclusive, as discussed below.

A study by Kane et al. examining the variation of solar variables in relation to the solar rotation period found the 10.7cm (2800Mhz) solar radio power flux spectra to exhibit peaks at 24.6, 18.2 and 14.3 days, although no degree of statistical certainty is attributed to the findings [*Kane, Depaula, and Trivedi*, 1995]. It is also difficult to view these figures as a good indicator of the periodic behaviour of the solar radio flux as the analysis was based on only 128 days of data, which is only 5 solar rotations, and is very small compared to the evolutionary timescale of the 11 year solar cycle .

Whilst there is some evidence from other studies for a near 27-day periodicity in the solar radio flux (e.g. [*Donnelly and Puga*, 1990]), the strong 27-day periodicity of Figure 8-5 could be due to the fact that the analysis is a superposed epoch analysis. Figure 8-5 was created by averaging superimposed extractions for 253 SPEs over an approximate 35 year span, synchronising each SPE occurrence at a point in time. In reality, active regions are known to occur simultaneously, giving rise to periodicities other than a simple 27-days. (For example, Mursula and Zieger examined the power spectra of a number of solar variables over a 20 year time span and concluded that 13.5 day periodicities were the result of two solar active regions approximately 180 degrees apart [*Mursula and Zieger*, 1996]). It is possible that significant 27-day periodicities exist only when emissions are dominated temporarily by a single active region lasting a few solar rotations, and this may explain why a 27-day periodicity was not dominant in Kane's study.

8.2.3 Performance Comparison of Predictor Variables

Examination of the radio and x-ray flux has given a useful visualisation of the predictor variable behaviour and has identified points in time that can offer optimal potential as inputs to the classification model. In this section, models were tested with different predictor variables, chosen based on their F-statistics.

For each predictor variable (XS, XL, XS/XL and the radio flux), 3-hour time intervals between –720 and –48 hours were ranked by the F-statistic. (24 intervals were chosen so as to be consistent with the models developed in the previous chapter, which took input vectors of 24 elements in the form of 3-hour averages over the 72-hour input window). Table 8-1 shows the top 24 time intervals for each predictor variable, ranked according to their corresponding F-statistic.

XS		**XL**		**XS/XL**		**Radio**	
Time/Hours	F-stat	Time/Hours	F-stat	Time/Hours	F-stat	Time/Days	F-stat
-120	*19.95*	-114	*19.22*	-120	*17.28*	-3	*27.69*
-96	*18.94*	-117	*17.49*	-126	*17.16*	-4	*22.58*
-123	*17.95*	-120	*17.33*	-96	*16.79*	-5	*22.35*
-117	*17.71*	-105	*17.27*	-78	*16.71*	-6	*18.21*
-105	*17.59*	-96	*17.02*	-123	*16.37*	-28	*16.80*
-126	*15.86*	-111	*16.28*	-102	*16.20*	-27	*15.84*
-102	*14.82*	-66	*15.69*	-150	*15.91*	-7	*14.88*
-111	*14.59*	-123	*15.53*	-171	*14.76*	-29	*13.98*
-78	*14.58*	-93	*15.29*	-54	*14.44*	-30	*13.90*
-150	*14.21*	-69	*14.72*	-168	*13.83*	-26	*13.16*
-75	*13.22*	-90	*13.94*	-117	*12.94*	-31	*12.47*
-54	*12.53*	-108	*13.61*	-606	*12.60*	-8	*11.82*
-168	*12.37*	-72	*13.44*	-174	*12.56*	-25	*10.99*
-114	*12.18*	-75	*13.09*	-105	*11.52*	-32	*10.23*
-93	*12.00*	-87	*11.48*	-75	*10.85*	-57	*9.23*
-99	*11.91*	-99	*10.68*	-99	*10.39*	-9	*9.17*
-72	*11.37*	-126	*10.62*	-222	*9.54*	-24	*8.70*
-108	*11.28*	-102	*10.37*	-144	*9.29*	-55	*8.20*
-69	*11.12*	-84	*10.33*	-324	*9.22*	-58	*8.19*
-171	*11.09*	-78	*9.78*	-111	*9.15*	-33	*8.01*
-66	*10.35*	-63	*9.58*	-147	*8.63*	-56	*7.82*
-144	*8.99*	-150	*8.93*	-702	*7.85*	-10	*6.68*
-51	*8.54*	-51	*8.76*	-654	*7.67*	-59	*6.39*
-174	*8.50*	-168	*8.64*	-720	*7.44*	-54	*6.34*

Table 8-1 Top 24 time intervals for each predictor variable ranked according to the F-statistic.

An MLP network classification model was created for each predictor variable in conjunction with a principal components analysis, and the first 6 principal components of the transformed predictor dataset were used as inputs to the network. Models were trained on 60 SPEs and 60 Quiet Periods and tested with the remaining 37SPEs and 113 Quiet Periods. In addition to creating a model for each of the predictor variables in Table 8-1, two further predictor sets were created. The first combined the XS and XL values by using the first 12 ranked time intervals of each of the XS and XL predictors; the second combined the XS/XL ratio and the solar radio flux by taking the top 12 ranked time intervals from each predictor. It can be noted that some of the inputs are actually from the time of the previous solar rotation (around –700 hours). Table 8-2 gives the average and standard deviation in performance of each MLP model configuration.

Inputs	Classification Success / % ±σ
Blanket XS/XL Ratio	62.6 ± 5.0
XS	61.9 ± 3.4
XL	59.8 ± 3.5
XS and XL	61.7 ± 2.9
XS/XL Ratio	64.7 ± 2.6
Radio Flux	58.8 ± 3.7
XS/XL Ratio and Radio Flux	63.3 ± 3.4

Table 8-2 Average performance of MLP classification models using different predictor variables as inputs. Uncertainty is one standard deviation.

It can be seen that all of the predictor variables generated very similar performance, with differences between models all being within one standard deviation.

The figure for the blanket XS/XL ratio denotes the MLP model from the previous chapter which used inputs from –120 to –48 hours. The performance of this model is only marginally below the performance average for the XS/XL inputs chosen via the F-statistics. This is probably because the blanket input vector already included the majority of time intervals with high F-statistics, and coupled with a principal components analysis, the negative effect of any 'poor' inputs on performance is also reduced.

The XS and XL predictors gave similar performance, but were slightly inferior to the x-ray ratio. This is a little surprising as the individual XS and XL channels generally exhibited higher F-statistics than the XS/XL ratio, as can be seen in Table 8-1. The ratio may be beneficial due to some form of coupling between inputs that is not readily visible from the graph but is exploitable by the network.

Use of the radio flux data as inputs generated the worst performance, which is surprising considering the relatively high F-statistics and seemingly large differences between the SPE and quiet period distributions in Figure 8-5. The reason for the comparatively poor performance must stem from some inseparability of the SPE and QP distributions. The F-statistic for the radio data is lent significance from the relatively large sample size, which reduces the uncertainty in the sample mean (i.e. the standard error) but examination of the standard deviation shows that there is still significant scatter about the SPE and quiet period means. Figure 8-10 plots the radio flux curves with error bars denoting standard deviation as opposed to standard error, where it can be seen that the distributions are no longer clearly separable.

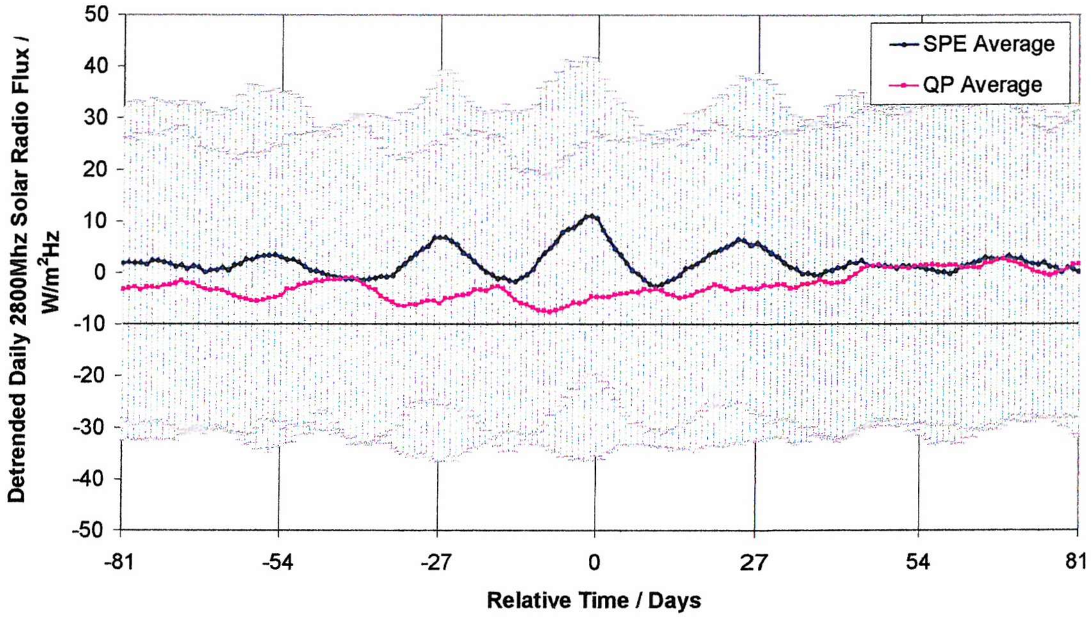

Figure 8-10 Variation in average detrended solar radio flux for SPEs and QPs for –81 days to +81 days relative to each event. Error bars denote standard deviation.

The same is true of the other predictor variables too, as is borne out by Figure 8-11 which shows the standard deviation as error bars on the SPE and quiet period cases for the XS/XL ratio.

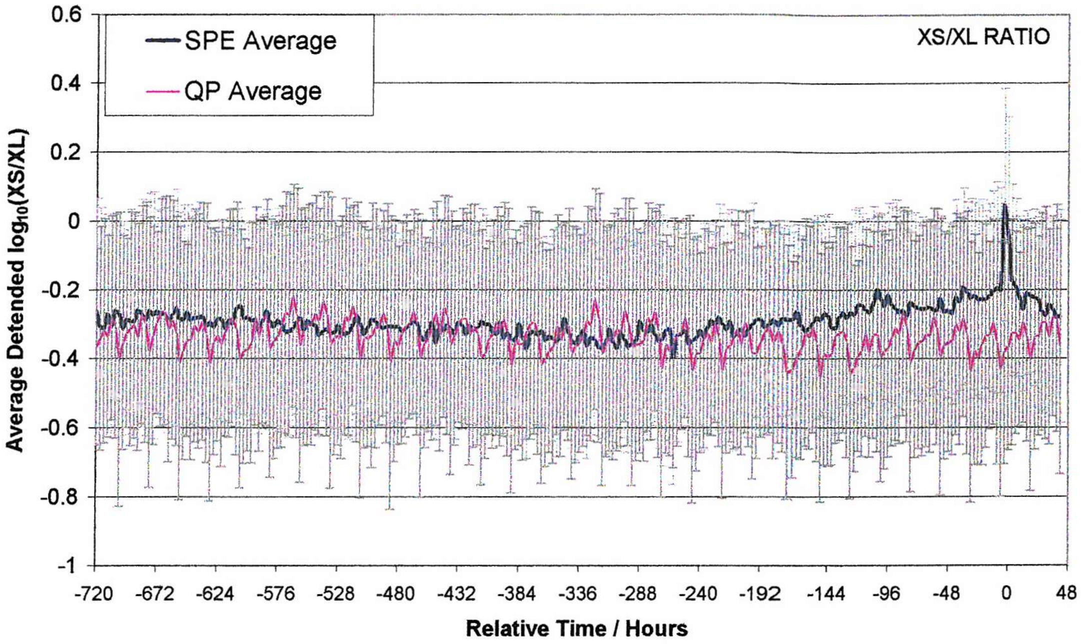

Figure 8-11 Variation in average XS/XL ratio as a function of time for SPEs and quiet periods. Error bars denote standard deviation.

Whilst the means of the distributions can be shown to be significantly different, there is still a high degree of scatter about the means, meaning that the distributions are not readily separable. This is undoubtedly the fundamental reason why classification ability is limited to the low 60% level: the inseparability is an inherent property of the dataset.

8.3 Summary

This chapter has visualised the behaviour of different predictor variables for the cases of SPEs and quiet periods and has performed a statistical analysis in order to quantify differences between the distributions. The effectiveness of different predictor variables has been compared by utilising them in MLP classification models, and the performance has been explained with respect to the properties of the predictor variable distributions.

Although GOES x-ray detectors measure flux from the full solar disk, this is almost totally composed of emissions from active regions. As the solar x-ray flux is not significantly attenuated by CMD angles of <90 degrees it means that GOES x-ray detectors are monitoring the variation in active region emissions over time.

Via the F-statistic it has been determined that the distributions for SPEs and quiet periods in XS, XL and the XS/XL ratio are significantly different in the mean to a greater than 99% certainty days before event occurrence. In the SPE case the average x-ray flux (and ratio) rises in comparison to the quiet period case from around 140 hours prior to time zero. A variation in solar x-rays over a timescale of days has not been previously documented in association with SPEs, and this is a new result.

Small peaks in the x-ray fluxes and ratio for the SPE case are present at around –720 hours, suggesting that the activity at time zero is recurrent and related to a developing active region. Evidence in solar x-rays for SPEs being related to recurring active regions has not been published elsewhere.

The average radio flux associated with SPEs exhibited a strong 27-day periodicity, passing through a maximum at the time of the event. This is the first time that SPEs have been associated with a 27-day periodicity in the solar radio flux, and builds on the findings of Chakravorti et al. who also found the solar radio flux to peak on the day of an SPE [*Chakravorti, Das, Sen, and Dasgupta*, 1991].

The 27-day recurrence in radio activity for the SPE case suggests that SPEs are related to active regions which develop over several solar rotations. This is supported by Mursula and Zieger, who found solar stream structures from active regions to last for around 4 solar rotations [*Mursula and Zieger*, 1996].

Analysis of the SPE dataset found 51% of SPEs to occur within 4.5 days of a local radio flux maximum, but the presence of any permanent 27-day periodicity in the solar radio flux has not been ascertained. Kane et al. did not find a 27-day periodicity to be dominant in samples of the solar radio flux, but this may have been due to small sample sizes [*Kane, Depaula, and Trivedi*, 1995]. Equally, the 27-day periodicity may only be present for a few solar rotations when SPE producing active regions are developing.

An examination of the 2800Mhz solar radio flux for the cases of SPEs and quiet periods has shown the SPE case average to differ significantly in the mean from the quiet period case up to two solar rotations prior to event occurrence. The average radio flux for quiet periods was found to be lower than that associated with SPEs, and was found to pass through a minimum at close to the time of the quiet period.

Selection of inputs using the F-statistics was found to give a small but not significant improvement to performance. A model using blanket inputs of the XS/XL ratio between –120 and –48 hours gave a 62.6% average classification success compared to 64.7% for XS/XL inputs chosen using the F-statistic. The blanket input model probably performs comparably because the majority of time intervals within the blanket input already have high F-statistics. Use of a principal components analysis has also minimised the negative contribution from any ‘poor’ inputs.

A comparison of predictor variables showed minimal difference in performance between input types. The XS/XL ratio was found to give the highest classification success of 64.7%, and radio data was found to generate the lowest classification success of 58.8%. Combining radio and XS/XL inputs was not found to be beneficial.

Performance is thought to be limited to around the low 60% level due to the inherent inseparability of the SPE and quiet period distributions. The distribution averages have been shown to be significantly different to a high level of confidence (lent by the large sample sizes), but scatter about the means is very high and generates significant overlap of the distributions.

Whilst performance of the classification technique appears to be limited to around the 60-65% level, such a degree of success is still an achievement considering a 48-hour lead time and the fact that a discrete x-ray flare is not needed as an input. A figure of >60% is still in line with recent ESA guidelines for SPE prediction requirements, and such a technique may still be useful for applications where the negative consequences of wrong predictions are small.

9. REAL TIME IMPLEMENTATION OF CLASSIFICATION MODEL

A main goal of the research activity was to produce a real time SPE prediction tool. This chapter describes the implementation of MLP networks within a real-time JAVA software framework and assesses their performance and behaviour over a 12-month period by using skill scores developed from terrestrial weather forecast models. The results and behaviour of the classification forecasting approach are explained further by considering the effects of solar x-ray flares, which are found to exhibit similar 'precursive' behaviour to SPEs.

9.1 The Real Time Model

A real time model was created by configuring the optimal MLP neural network model developed in chapter 7 to use real-time data from the GOES satellites. The model took 3-hour averages of the $\log_{10}$ detrended GOES XS/XL ratio over a 72-hour period to form an input vector. The first 6 principal components of the transformed input vector where than taken as inputs to the network in order to make forecasts with a 48 hour lead time. Earlier testing showed this model to have a 62.6% classification success rate.

Although in chapter 7 RBF models were found to have slightly superior performance to MLP models they were part of a larger analysis package called TSAR, and the RBF networks could not be extracted as stand alone objects. By contrast, MLP models developed in Neuframe could be extracted as raw code and offered a faster and easier way of incorporating models into a real-time framework. The performance difference between MLP and RBF models was shown to be of the order of 3%, which is within one standard deviation of the natural variation. The choice of model type is therefore not thought to be significant.

Models using inputs selected using the F-statistic in Chapter 8 were examined after the real-time model had been developed, and hence were not available for consideration. (However, it was shown in Chapter 8 that models using the XS/XL ratio chosen using the F-statistic were only 2% better than the blanket ratio input, therefore any performance gain from using the F-statistic in a real-time model will again not be significant).

9.1.1 Software Architecture

One of the principal requirements of the model was the retrieval of remote data in real time. This led to the software being written in JAVA due to its superior ability to handle internet connections. JAVA also has the benefit of being non-platform specific.

The software was denoted 'Predictor of Proton Events' (POPE) and was written as a stand-alone program in accordance with ESA software standards for small projects*. POPE was designed so that prediction models could be 'bolted' to a real-time retrieval framework with minimal effort, therefore facilitating the implementation of any other prediction models at a future stage. Compliance with ESA software standards necessitated the creation of a software specification document and a user requirements manual. These can be found in Appendices G and H and give a detailed presentation of the software architecture and specific code modules. A schematic diagram of the software process is given in Figure 9-1.

The prediction model within POPE consisted of 10 MLP models acting in parallel to form a multiple model configuration. The distinct tasks carried out by POPE can be broken down as follows:

- Retrieval of real-time GOES satellite data from a remote ftp site.
- Creation of a continuous data file made by stitching together several GOES data files.
- Detrending of x-ray data to remove the long term solar cycle trend.
- Creation of a delay vector from continuous GOES x-ray data.
- Pre-processing of the delay vector (normalising, PCA and scaling)
- Running a neural network with the processed delay vector as an input.
- Interpreting and recording the network outputs and other information relevant to each prediction.

---

* ESA Software Standard PSS-05

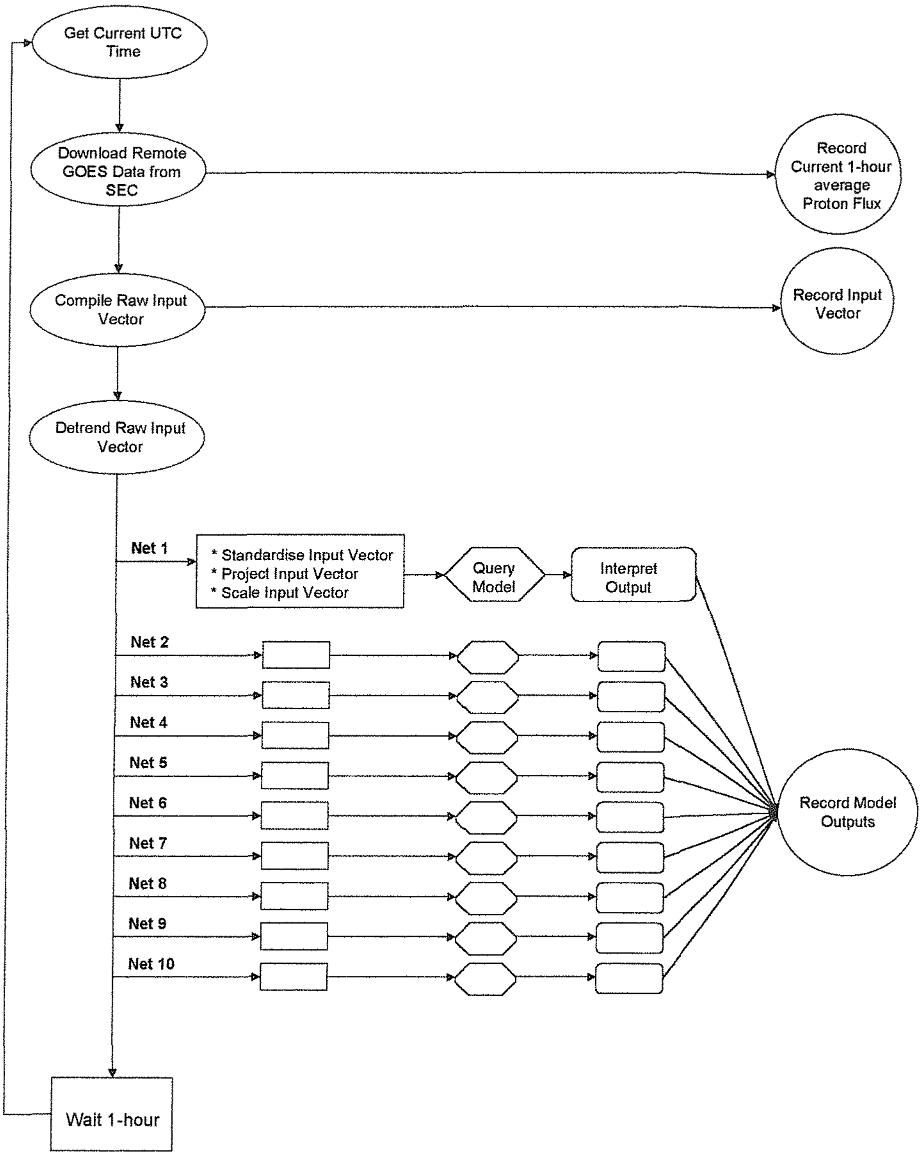

Figure 9-1 Schematic diagram of real time model process.

POPE places the process within a control loop which repeats every 60 minutes, creating predictions along a rolling time line at the rate of 1 per hour. The repeating interval was chosen so as to be smaller than the resolution to which the start times of SPEs were found in the training dataset. (The start times of SPEs were found to the nearest 2-hours).

Each time the model runs the output file records the numerical output from each of the 10 component MLP networks and, using a decision boundary of '50', records the number of networks from 10 that are predicting an SPE 48-hours from the run time. A thorough description of the POPE output can be found in the user manual in the appendix.

POPE was put into real-time operation during December 2001 utilising x-ray ratio data from the GOES-8 satellite as inputs. Since then approximately 12-months of output data has been generated, allowing an initial assessment of operational performance to be made.

9.2 Performance of the Real Time Model

The predictions made by POPE between 3/12/01 and 16/12/02 have been assessed by breaking down the observed proton flux time series into a binary time series of SPES and Quiet Periods with 1-hour resolution. The observed series has then been compared against the POPE forecast. Within the training set the start times of SPEs had been found to the nearest two hours, hence the predictions from the POPE model had a lead time of 48-hours, ± 2 hours.

9.2.1 Method

The output from POPE that was used to assess its performance consisted of an integer of between 0 and 10, representing the number of networks from 10 which had an output of >50, i.e. the number of networks from 10 which were predicting an SPE in 48-hours time. This is a 'majority vote' output, which was shown previously to be superior to a single network.

Over the period of operation 1-hour resolution GOES-8 proton data was examined and an SPE was defined as occurring if the hourly averaged >10MeV integrated proton flux was greater than 1.0 p.f.u. for 12 consecutive hours. This is consistent with the way in which SPEs were defined in the original training set. The start time of each SPE was taken as the time at which the hourly averaged proton flux was first above 1.0 p.f.u. SPEs were identified as ending when the hourly averaged >10MeV integral proton flux fell below 1.0 p.f.u. for 12 consecutive hours. The POPE prediction for each SPE was taken as the highest output from the model from within ± 2 hours of the SPE start time. Data collected during an ongoing SPE was ignored, thus the POPE model was assessed on its ability to predict the initial onset of an SPE.

Observed values with a >10MeV hourly averaged proton flux of less than 1.0 p.f.u. which occurred at least 6 hours after the end of an SPE were classed as instances of quiet periods. The prediction for each quiet period was taken as the highest output from the model from within ± 2 hours of the quiet period.

Outputs from the POPE model were only utilised if they were based on an input delay vector that did not contain a significant amount of missing data (due to data gaps from GOES). An input vector was deemed unsuitable if any of its component 3-hour averages contained more than 50% missing data.

The 48-hour forecast from the model was matched to each observed value of the hourly averaged >10MeV integrated proton flux. The forecast consisted of an integer of between 0 and 10 and the threshold used to interpret the forecast was altered in order to examine the effect on SPE and quiet period classification success rates.

Statistic		Definition	
Prob. of Detection 'Yes'	PODy	$a/(a+c)$	Proportion of 'Yes' observations that were correctly forecast
Prob. of Detection 'No'	PODn	$d/(d+b)$	Proportion of 'No' observations that were correctly forecast
False Alarm Ratio	FAR	$b(a+b)$	Proportion of 'Yes' forecasts that were incorrect
Bias		$(a+b)/(a+c)$	Ratio of number of 'Yes' forecasts to number of 'Yes' observations
Critical success index	CSI	$a(a+b+c)$	Proportion of hits either forecast or observed.
True Skill Statistic	TSS	PODy+PODn-1	Measures ability to discriminate between 'Yes' and 'No'.
Heidke skill score	HSS	$(a+d-C_1)/(N-C_1)$	Percent correct, corrected by those expected correct by chance
Gilbert skill score	GSS	$(a-C_2)/(a+b+c-C_2)$	CSI, corrected by number of hits expected by chance(C_2).
Where:	N	$A+b+c+d$	Total number of events
	C_1	$C_2+(b+d)(c+d)/N$	Number expected correct by chance
	C_2	$(a+c)(a+b)/N$	Number of hits expected by chance
		a	Number of Hits ('Yes' forecast and 'Yes' observed)
		b	False Alarms ('Yes' forecast but not observed).
		c	Misses ('No' Forecast and 'Yes' observed)
		d	Correct Nulls ('No' Forecast and 'No' observed)

Table 9-1 Description of statistics utilised to compute skill scores. (Taken from [*Smith, Dryer, Ort, and Murtagh*, 2000])

The POPE model has been compared to the SEC 2-day forecast model over exactly the same operational period. The SEC 2-day forecast is based on human judgement and experience and consists of a probability of SPE occurrence in two days. In order to compare the models statistical skill scores have been calculated based on methods presented by Smith et al. who derive skill scores for an interplanetary shock prediction model [*Smith, Dryer, Ort, and Murtagh*, 2000]. Skill scores are summarised in Table 9-1.

9.2.2 Results

Since the commencement of real-time operation in December 2001 13 SPEs were found to have occurred. These are shown in Table 9-2 along with their observed fluence and the highest output from POPE within ±2 hours of the SPE start time.

ID	Start Time	Duration / Hours	Fluence / p/cm^2	POPE Output
1	26/12/01 6:00	49	3.54E+08	10
2	29/12/01 5:00	17	2.35E+07	10
3	30/12/01 21:00	119	2.33E+08	8
4	10/01/02 20:00	67	1.03E+08	9
5	18/03/02 13:00	32	3.23E+07	2
6	21/04/02 2:00	113	2.84E+09	6
7	22/05/02 18:00	44	1.05E+08	5
8	16/07/02 18:00	46	1.02E+08	1
9	22/07/02 7:00	97	8.31E+07	Model Down*
10	22/08/02 4:00	20	1.65E+07	Model Down*
11	24/08/02 1:00	60	3.13E+08	10
12	07/09/02 7:00	19	2.55E+07	10
13	09/11/02 19:00	34	1.42E+08	9

***Due to internet server losing connection to GOES data site.**

Table 9-2 SPEs identified during the period of real-time operation.

For two SPEs the POPE model was not running, but for the remaining 11 SPEs the POPE output was at least 6 in all but three cases. Using a majority vote of 5 or more to indicate an SPE, the POPE model correctly predicts 9 of the 11 SPEs. This indicates excellent SPE prediction ability, but a true measure of performance can only be made by considering the success rate for quiet period detection as well as SPE detection.

6249 quiet periods were observed during the POPE operational periods and Table 9-3 shows the classification success for each category for different threshold interpretations of the POPE output. (The POPE output is an integer from 1-10 indicating the number of networks that are predicting an SPE).

Threshold	>5for SPE	>6 for SPE	>7 for SPE	>8 for SPE	>9 for SPE
SPEs from 11	8	7	7	6	4
QPs from 6249	2158	2711	3345	4120	5123
SPEs / %	72.7	63.6	63.6	54.5	36.4
QPs / %	34.5	43.4	53.5	65.9	82.0
Average / %	**53.6**	**53.5**	**58.6**	**60.2**	**59.2**

Table 9-3 Classification success for real-time POPE model for different threshold interpretations of the multiple model output. Average success is shown in bold.

A >5 threshold generates a high SPE classification rate (72.7%) but only identifies 34.5% of quiet periods correctly, indicating that quiet periods tend to be misclassified as SPEs by the model. Optimal overall success of 60.2% is achieved when the threshold is set to >8, i.e. more than 8 of the 10 networks need to be predicting an SPE before the overall model output can be interpreted as an SPE forecast. The fact that a high threshold generates optimal average success indicates that the model has generally over predicted, meaning that the model has a tendency to generate high outputs for Quiet Periods.

Varying the threshold used to interpret the multiple model output can act as a way to bias the output to specific user needs. If the consequences of a false alarm can be tolerated a lower threshold can be adopted to miss fewer SPEs, whereas if a low false alarm ratio is important a high threshold can be implemented with the risk of missing more SPEs. In real terms quiet periods occur far more frequently than SPEs. Observations during the trial period indicate a ratio of 1:700 for their respective occurrence rates, hence a model that is biased towards correctly identifying quiet periods will be correct a larger proportion of times.

The POPE model was based on 10 MLP neural networks which had an average classification success of 62.6% when exposed to the test set during development. Operational success is 60.2%, indicating that the models have performed comparably in real-time operation. This validates the success rates that were obtained during testing, although it must be noted that the sample size of SPEs in real-time operation is small and may not be an accurate measure of SPE prediction ability.

Table 9-4 shows a contingency table for the POPE classification model using >8 as the decision threshold. The SEC 2-day forecast was examined over exactly the same operational period as POPE, and by categorising probabilities of ≥0.5 as an 'SPE' prediction and probabilities of <0.5 as quiet periods a direct comparison has been made between the two techniques. Table 9-5 is a contingency table for the SEC 2-day forecast. For completeness a contingency table for the PROTONS model is given in Table 9-6 but pertains to performance during the calendar year of 1989.

POPE Model		Event Observed	
		Yes	No
Event Forecast	Yes	6	2129
	No	5	4120

Table 9-4 Contingency table for real-time operation of POPE model, Dec2001-Dec2002.

SEC 2-Day Forecast		Event Observed	
		Yes	No
Event Forecast	Yes	1	4
	No	12	332

Table 9-5 Contingency table for real-time operation of SEC 2-day SPE forecast , Dec2001-Dec2002.

PROTONS Model		Event Observed	
		Yes	No
Event Forecast	Yes	17	16
	No	5	63

Table 9-6 Contingency table for real-time operation of SEC PROTONS model during 1989 [*Heckman, Kunches, and Allen*, 1991].

Table 9-1 summarises the statistical formulas as presented by Smith et al. for the appraisal of interplanetary shock prediction models based on a 2-category forecast [*Smith, Dryer, Ort, and Murtagh*, 2000]. The statistical skill scores have been calculated for POPE based on the contingency table of Table 9-4 and are shown in Table 9-7. Skill scores have also been calculated for the PROTONS model and the SEC 2-Day forecast based on their respective contingency tables. The skill statistics serve as a formal method for model comparison, and can also be used to infer the behaviour of a model.

Statistic		POPE	SEC 2-Day	PROTONS
Number of Hits ('Yes' forecast and 'Yes' observed)	a	6	1	17
False Alarms ('Yes' forecast but not observed).	b	2129	4	16
Misses ('No' Forecast and 'Yes' observed)	c	5	12	5
Correct Nulls ('No' Forecast and 'No' observed)	d	4120	332	63
Prob. of Detection 'Yes'	PODy	0.545	0.077	0.773
Prob. of Detection 'No'	PODn	0.659	0.988	0.797
False Alarm Ratio	FAR	0.997	0.800	0.485
Bias		194.091	0.385	1.500
Critical success index	CSI	0.003	0.059	0.447
True Skill Statistic	TSS	0.205	0.065	0.570
Heidke skill score	HSS	0.002	0.092	0.483
Gilbert skill score	GSS	0.001	0.048	0.318
Total number of events	N	6260	349	101
Number expected correct by chance	C_1	4121.503	331.372	60.376
Number of hits expected by chance	C_2	3.752	0.186	7.188

Table 9-7 Skill score statistics for POPE and PROTONS calculated from contingency tables.

The skill scores for POPE in terms of the Critical Success Index (CSI), Heidke Skill Score (HSS) and Gilbert Skill Score (GSS) are all close to zero, which indicates poor performance, and indeed these values are lower than the same scores from the 2-day SEC forecast model.

The reason why the POPE model generates low skill scores is predominantly due to the regime in which it operates. In the case of the POPE model SPEs account for only 0.14% of the total observations, meaning that it only takes a small percentage of Quiet Periods to be misclassified to form a large proportion of the total SPE forecasts. For example, even if 99.9% of the 6249 Quiet Periods are forecast correctly, 6 Quiet Periods are still misclassified as being SPEs, which equates to 50% of the total number of SPE predictions. In fact, Table 9-7 shows a false alarm ratio of 0.997 for POPE, which shows that 99.7% of SPE predictions were false alarms. The large imbalance between the number of observed SPEs and observed Quiet Periods means that the skill scores for POPE are highly sensitive to false alarms. The imbalance in SPE and QP outcomes is a property of the dataset though, therefore any model operating in such a regime is inherently likely to produce a large number of false alarms.

The SEC 2-day forecast is the most directly comparable to the POPE model in that it predicts the likelihood of an SPE for the same lead time. The CSI, HSS and GSS skill scores for the SEC 2-day forecast are roughly an order of magnitude greater than for POPE, apparently indicating superior performance, but the difference in skill scores can in fact be attributed to the difference in operating regimes between the models. The POPE model makes 24 forecasts per day, whereas the SEC forecast is issued only once per day. The result is that the POPE model must predict far more quiet periods (an order of magnitude more) than the SEC forecast for the same number of observed SPEs, hence POPE is inherently more likely to generate false alarms and thus generate lower skill scores. The order of magnitude difference in skill scores can be directly attributed to the order of magnitude difference between the QP:SPE ratio in the observed outcomes for each model.

A fairer measure of 'skill' is given by the True Skill Score (TSS), which measures the ability of the model to discriminate between the two outcomes irrespective of the relative frequency of the outcomes in the observed distribution. It is not weighted by the total number of correct predictions, but by the percentage of each category that is predicted correctly. It can be seen from Table 9-7 that POPE scores approximately 3 times higher than the SEC 2-day forecast, indicating that it is in fact better at distinguishing between SPEs and Quiet Periods. In effect the SEC 2-day forecast predicts a Quiet Period in nearly all cases (as can be seen from its contingency table) and only predicted 1 of 13 SPEs during the trial period, hence it can never be regarded as a reliable predictor of SPEs. However, because nearly all days are Quiet Periods it actually means that the SEC 2-day forecast is correct far more often than the POPE model, and this is why it scores well in the other skill statistics. In practice, the POPE model would be best suited to applications in which it is more important to predict SPEs than quiet periods, i.e. applications in which false alarms can be tolerated.

The PROTONS model is not directly comparable to the POPE model because it has an effective lead time of just 0-6 hours to SPE onset, but it does represent the best available forecast tool for SPEs at the moment in terms of prediction accuracy. The skill scores for PROTONS are higher than for POPE and the SEC 2-day forecast, although this can in part be attributed once again to a difference in operating regimes. The PROTONS model is run only after a significant x-ray flare has occurred, hence it is only run at times when there is an increased likelihood of an SPE occurring (and hence a reduced likelihood of no SPE occurring) and it can be seen from the contingency tables that SPE occurrences make up 28% of observations for PROTONS, compared to just 0.14% for POPE. The more equal observed ratio of the two outcomes means that the PROTONS model is less prone to generating false alarms, and this is borne out by the false alarm ratio in Table 9-7, but the relatively good performance cannot be entirely attributed to the operating regime. The True Skill Score for PROTONS is still over twice that of POPE and shows that the technique is at least 'semi-intelligent' rather than just playing to the statistics of the observed distribution like the SEC 2-day forecast.

The POPE model is better at predicting SPEs than the current SEC 2-day forecast, but the PROTONS model still represents the best model in terms of forecast accuracy.

The POPE model needs to be improved by reducing the number of quiet periods that are classified as SPEs whilst at the same time maintaining a high SPE detection rate. One method of doing this is to change the regime in which the model operates in order to better balance the ratio of observed QPs to SPEs: a lower frequency of predictions would result in fewer observed quiet periods for the same number of observed events.

However, given that the current model has been trained using input vectors from exactly 48-hours±2hours prior to event occurrence it is difficult to justify a forecast repeat time that is greater than the resolution to which the SPE start times were found. Predictions made by POPE are effectively only valid for a 4 hour period, hence the minimum run-frequency is 6 forecasts per day for continuous coverage of the timeline. Assuming that a 4 hour repeat cycle is adopted this would quarter the total number of forecasts from 24 per day to just 6, but even so, based on contingency Table 9-4, around 500 misclassified Quiet Periods would still occur, which is over 80 times the number of observed SPEs, and would still result in a very high false alarm ratio. Rather than trying to alter the operating regime of POPE it is perhaps better to try and understand why the model produces so many false alarms in the first place.

One advantage that POPE does have over existing forecasting techniques (both the 2-day forecast and PROTONS) is that it is completely autonomous and requires no supervision, 'decision to run' or other human input. The POPE model is not susceptible to the subjectivity of human judgement and can be left to run in the background with no requirement for an operator. This means that POPE has attractively low resource requirements and is easy to install and integrate within a space weather forecasting service.

9.3 Summary

This chapter has described the implementation of an MLP neural network model into a real-time framework and has assessed its performance as an operational SPE forecasting tool. The model was designated 'Predictor of Proton Events' (POPE) and skill score statistics were used to compare the model against existing SPE forecasts made by the SEC.

During a 12-month period between December 2001 and December 2002 11 SPEs occurred whilst the POPE model was operational. 6 of the 11 events (55%) were successfully predicted, as were 4120 of 6249 quiet periods (66%).

The POPE model was compared to the SEC 2-day forecast over exactly the same operational period and was found to have a True Skill Score over 3 times greater, indicating a better ability to distinguish between SPEs and Quiet Periods. The POPE model is an improvement over the existing 2-day lead time forecast for SPEs in that it has superior discrimination.

The current 2-day SPE forecasts made by the SEC are strongly biased to Non-SPE prediction with virtually no SPE prediction capability, with the result that 1 of 13 SPEs (8%) and 332 of 336 quiet periods (99%) were forecast correctly over the operational period. The SEC 2-day forecast is nearly always a Quiet Period, hence is correct in nearly all cases, but the model has very little ability at forecasting an SPE.

Despite a better true skill score, the POPE model output was found to be wrong more often than the 2-day SEC model by virtue of the fact that significantly more quiet periods than SPEs are present in the observed proton flux timeline. This observed distribution led the number of misclassified quiet periods to be very large in comparison to the number of correctly forecast SPEs and produced a very high false alarm ratio which governed the skill score statistics (other than the True Skill Score).

The SEC PROTONS model still offers the highest accuracy for SPE forecasting with forecast success rates of 77% for SPEs and 80% for Quiet Periods respectively, but achieves this by trading lead time and imposing the pre-requisite of a solar x-ray flare. This compares to success rates of 55% and 66% respectively for 2-day forecasts made by POPE.

The performance of the POPE model suffered due to the regime in which it operated. It was identified that the POPE model could be improved by reducing the number of false alarms as a fraction of the total number of SPE predictions.

The POPE model in its current form still offers significant potential as a real-time operational forecast tool and is superior to the current 2-day SEC SPE forecast. One significant advantage of the POPE model is the benefit of being able to operate autonomously without the need for human input or supervision, which could be favourable in some applications.

10. CONSIDERATION OF X-RAY FLARES

SPEs are known to have a strong association with solar x-ray flares, and this is borne out by the earlier result from chapter 6 in which 91 of 98 SPEs could be temporally matched to x-ray flare occurrences.

This high association between SPEs and solar x-ray flares means that when x-ray data was extracted based on the start times of SPEs, it was also, by proxy, extracted about the time of a solar x-ray flare. No consideration of x-ray activity was made in defining the quiet periods within the training set, but it is probable that no x-ray flares were present at time zero in the quiet period case. Consequently this means that during training the neural networks may have been learning to distinguish between the cases of 'x-ray flare' and 'no x-ray flare' as opposed to 'SPE' and 'non-SPE'. If this is the case then it may explain the high false alarm rate of the real-time POPE model. The model may have been triggered into giving SPE predictions by the relatively common occurrence of solar x-ray flares.

In order to test this hypothesis the POPE model has been tested with a large dataset of query inputs pertaining to times at which an x-ray flare occurred at time 0, but no proton event occurred. This has been contrasted with the model's response to an equivalent dataset of quiet periods which were selected at random with no consideration of x-ray activity. The behaviour of solar x-rays prior to x-ray flares has also been studied and is compared to the behaviour of solar x-rays prior to SPEs.

To try and account for the occurrence of x-ray flares the POPE models were re-trained using quiet periods taken at the times of non-proton x-ray flares. In addition, the abundance of quiet periods in relation to SPEs in the observed data was reproduced within the training set to see if this improved the ability to predict quiet periods correctly.

10.1 X-ray flare dataset

As described in section 6.4 flare listings from the SEC allowed flare associations to be made for all SPEs occurring after 1986. These flare listings identified times at which x-ray flares occurred but with no associated proton event.

Considering solar active years between 1986 and 1999, a total of 17731 x-ray flares of class X, M and C were recorded. 91 of these flares could be associated with the onset of an SPE and were classified as being proton associated. A further 4840 flares were found to occur during an SPE, or within 3 days of the end of an SPE, and these were removed from the dataset. The remaining 12800 flares were designated 'quiet flares' as they were associated with background levels of solar proton flux. The breakdown of the flares is given in Table 10-1 and Table 10-2.

	SPE associated Flares		Quiet flares	
X-Ray Flare Class	Number	% of Sample	Number	% of Sample
C	10	10.2	11400	89.1
M	46	46.9	1327	10.4
X	35	35.7	73	0.6
No flare	7	7.1	-	-
Total	98		12800	

Table 10-1 Solar x-ray flares with SPE associations and no SPE associations occurring in solar active years between 1986 and 1999.

X-Ray Flare Class	Percentage of all Flares	Percentage of Flares associated with SPEs	Percentage of SPEs accounted for
C	88.5%	0.1%	10.2%
M	10.7%	3.4%	46.9%
X	0.8%	32.4%	35.7%
All	100.0%	0.7%	92.9%

Table 10-2 Proportional breakdown of solar x-ray flares occurring in solar active years between 1986 and 1999 with the proportion of SPEs that can be associated with them.

The statistical breakdown of the sample shows that overall, whilst 92.9% of SPEs can be associated with x-ray flares, only 0.7% of x-ray flares can actually be associated with SPEs. The correlation is clearly not 1:1 hence SPEs cannot be predicted simply by successfully predicting a flare.

However, it is interesting to note as an aside that whilst X class flares constitute only 0.8% of all flares, 32.4% of these flares can be associated with SPEs. This suggests that if X-ray flares could be forecast successfully it would follow that there is a 32.4% chance of an SPE also occurring.

10.2 Model response to x-ray flares

10.2.1 Method

Using the x-ray flare listings from the SEC a test set of 5843 quiet periods was constructed, consisting of detrended x-ray ratio extractions with a 'quiet flare' at time 0, i.e. an x-ray flare occurred at time 0 with no associated proton event. The proportion of flares within the test dataset was kept approximately equal to the sample proportions of Table 10-1 and is given below in Table 10-3.

Quiet Flare Type	Number of examples
C	5184
M	624
X	35
total	5843

Table 10-3 Dataset of quiet flares for which predictor data was extracted.

A further 3000 examples were generated by selecting times at random between 1986 and 1999 and extracting detrended x-ray ratio values. These examples represent times at which no attempt has been made to place a solar x-ray flare at time 0 in the extraction window. The POPE model was queried with the 'quiet flare' dataset and the 'random quiet period' dataset and the classification success rates were compared.

10.2.2 Results

The POPE output was interpreted using a threshold of >8, this being previously identified as optimal in section 9.2.2, i.e. as long as 2 or more networks were indicating a quiet period, the overall response from POPE was taken as a quiet period. Results are given in Table 10-4.

	Quiet Flares Successfully Forecast as QPs	Random Quiet Periods Successfully Forecast as QPs
X Flares	60.0%	-
M Flares	59.1%	-
C Flares	66.1%	-
Total	**61.8%**	**90.5%**

Table 10-4 POPE performance with test set of 5843 quiet flares and 3000 random quiet periods.

A comparison of success rates indicates that quiet periods that coincide with flares are often misclassified as SPEs whereas times at which there is no x-ray flare are successfully predicted as quiet periods.

When examining the classification rates by flare type there is little difference in success rates, indicating that even small flares cause the model to forecast an SPE. For the randomly selected quiet periods it can be assumed that there is generally no flare at time zero, and in these cases classification success is very high at >90%. The result indicates that it is highly likely that the large number of false alarms produced by the POPE model are caused by x-ray flare precursors being wrongly identified as SPE precursors.

10.3 Retraining with quiet flares

A possible solution to x-ray flares being falsely interpreted as SPEs is to re-train the POPE models with quiet periods specifically taken at the times of x-ray flares in the hope that the network can learn to differentiate between 'quiet flares' and x-ray flares associated with SPEs.

10.3.1 Method

A dataset was constructed using 98 SPEs and 660 'quiet flares' occurring in solar active years between 1986 and 1999. The ratio of flare categories within the quiet flare dataset was kept approximately equal to that in the observed sample of Table 10-2. The dataset was divided into a training and a test set in which the ratio of SPEs to quiet periods was kept at 1:1 in the training set. A breakdown of the dataset set is given below in Table 10-5.

	SPE Dataset			Quiet Flare Dataset		
	Train	Query	Total	Train	Query	Total
X flare association	21	14	35	1	10	11
M Flare association	28	18	46	7	70	77
C Flare association	6	4	10	52	520	572
No Flare association	4	3	7	-	-	-
Total	59	39	98	60	600	660

Table 10-5 Dataset proportions on which the POPE networks were retrained.

10 permutations of the training and query sets were used to retrain and test the 10 component MLP networks of POPE.

10.3.2 Results

The 10 retrained networks were assessed individually and the average classification success was recorded. This is compared against the classification success when randomly selected quiet periods were used (i.e. the result from section 7.2.8) in Table 10-6 below.

	Using quiet flares in training set	Using random quiet periods in training set
SPE Classification Success	66.4%	69.5%
Quiet Period classification success	34.9%	55.8%
Average Overall classification success	**50.6%**	**62.6%**

Table 10-6 Comparison of performance between POPE MLP models when trained with randomly selected quiet periods and when trained with 'quiet flare' periods.

It can be seen that when the MLP networks are trained using quiet periods that have x-ray flare associations the classification success drops to 50%, which is no different to the performance expected from a random output.

The result indicates that the network cannot learn to differentiate between x-ray flares with proton associations and x-ray flares which have **no** proton associations. This suggests that the behaviour of the predictor variable prior to x-ray flares is the same regardless of whether or not the flare is associated with an SPE.

10.4 Consideration of quiet period abundance

Another possible way to try and reduce the false alarm rate is to condition the model to quiet periods by using a training set that contains representative proportions of SPEs and quiet periods. By running on a rolling time line at a rate of 1 prediction per hour it has been shown that only 0.14% of the observations seen by POPE are SPEs hence more than 99% of occurrences are quiet periods. It may be beneficial to increase the ratio of quiet periods to events in the training set in order to reproduce the observed ratio more closely.

10.4.1 Method

The 98 SPEs from Table 10-5 were divided to produce a training set of 59 examples, necessitating approximately 3000 quiet periods in order to form a training set composed of 2% events. (This is close to the observed ratio of 1%, but could not be reproduced exactly as it would have required an impractically large data extraction). 2922 quiet periods with flare associations were placed with the 59 SPEs to form the training dataset, with a further 2921 quiet flare periods serving as test data. A summary of the dataset is given below in Table 10-7.

	SPE Dataset			Quiet Period Dataset		
	Train	Query	Total	Train	Query	Total
X flare association	21	14	35	18	17	35
M Flare association	28	18	46	312	312	624
C Flare association	6	4	10	2592	2592	5184
No Flare association	4	3	7	-	-	-
Total	59	39	98	2922	2921	5184

Table 10-7 Summary of the dataset in which the proportion of SPEs in the training set was set at approximately 2%.

The 10 component MLP networks within POPE were retrained and tested using 10 different random permutations of the dataset shown above.

10.4.2 Results

The average classification success of the test set was calculated and is summarised below in Table 10-8.

SPE: QP training ratio of 2:100	
SPE Classification Success	0.0%
Quiet Period classification success	100.0%
Average Overall classification success	**50.0%**

Table 10-8 Average classification performance of POPE component networks when trained using a training set of 2% SPE examples.

It is obvious from the result that increasing the proportion of quiet periods in the dataset does not allow the model to learn effectively. The significantly larger proportion of quiet periods means that minimum training error is produced if the model simply produces '0' for each example, as this will correctly identify 98% of the training set.

A review of the numerical outputs from the networks indicated that '0' was indeed produced for every example in the test set during the query phase. In real-time operation such a model would actually be correct most of the time, simply because quiet periods occur most of the time, but such an approach is clearly not intelligent.

10.5 Behaviour of predictor variables prior to x-ray flares

The previous results show that the classification technique cannot adequately distinguish between cases of quiet flares and proton associated flares. In order to quantify the behaviour of the predictor variables prior to x-ray flares the XS/XL ratio has been plotted over the input window for the cases of quiet x-ray flares and compared to the case of SPEs and to the case of randomly selected quiet periods.

10.5.1 Method

The average XS/XL ratio over the period –120 hours to –48hours was calculated for the dataset of quiet x-ray flares shown in Table 10-3. This is compared to the average behaviour of the SPE case (computed from the 98 SPEs in Table 10-1) and the average behaviour of 3000 randomly selected quiet periods (i.e. quiet periods which were not specifically based around the time of an x-ray flare). In addition, the detrended solar radio flux was also extracted and plotted for the same cases over a period spanning –81 to +81 days relative to time zero of the input window.

10.5.2 Results

Figure 10-1 below compares the variation in the average XS/XL for the cases of SPEs, randomly selected quiet periods and quiet periods associated with an x-ray flare (of class X, M or C) at time 0.

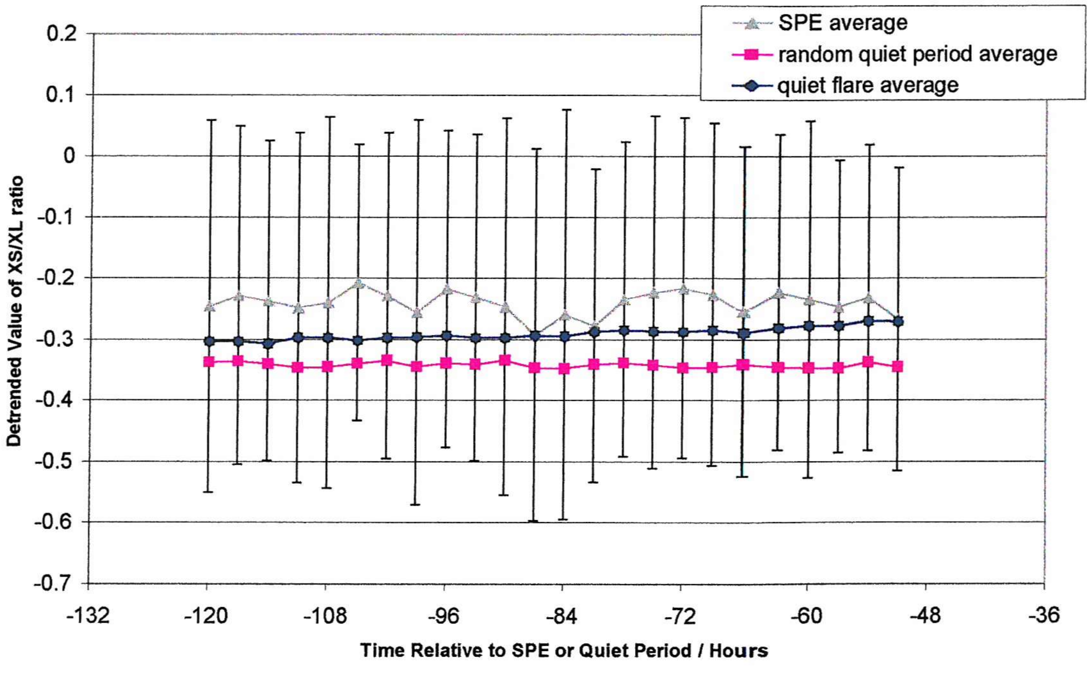

Figure 10-1 Comparison of XS/XL ratio over the input window for the cases of SPE, quiet flare periods and randomly selected quiet periods. Error bars denote one standard deviation and for clarity are shown only for the SPE case.

It is clear from Figure 10-1 that the quiet flare average lies between the random quiet period average and the SPE average, thus it is bound to represent a classification problem for a model which has only been trained on cases of SPEs and random quiet periods. Figure 10-2 shows the same comparison in the 2800Mhx solar radio flux.

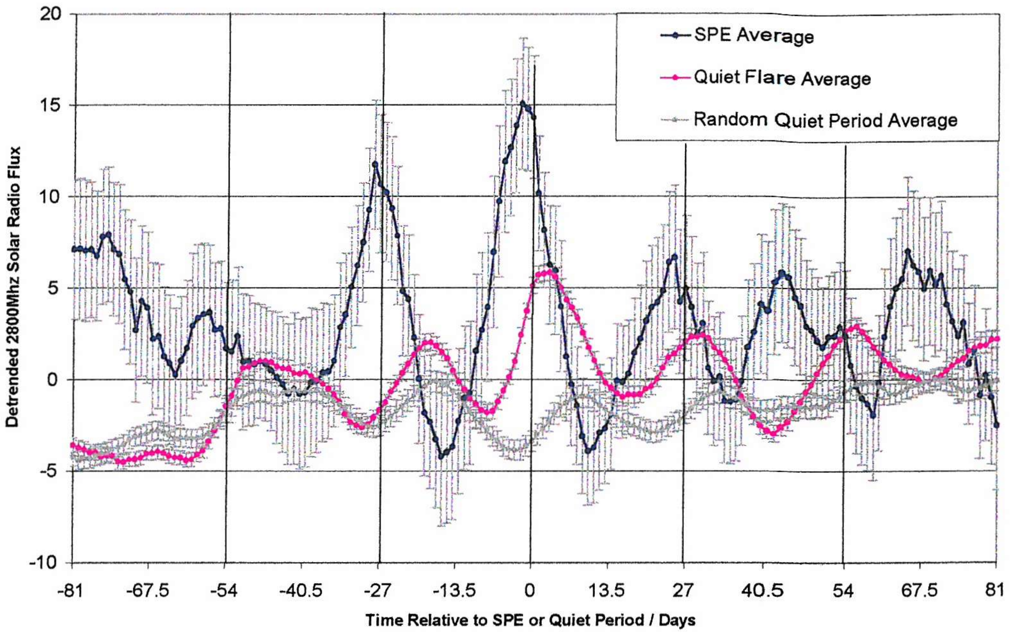

Figure 10-2 Comparison of detrended 2800Mhz radio flux over ±81 days relative to the time of event for the cases of SPE, quiet flare, and randomly selected quiet periods. Error bars denote standard error.

Figure 10-2 shows a similar situation with the solar radio flux in that the quiet flare case is somewhere between the SPE and than the randomly selected quiet period case. This again shows that the behaviour of solar variables prior to x-ray flares is similar to that prior to SPEs. It follows that future development may benefit from the use of three training classes, corresponding to cases of 'Flare with SPE', 'Flare with no SPE' and 'Quiet' (i.e. no flare and no SPE).

As an aside, Figure 10-3 shows the behaviour of the average solar radio flux classified by flare type for the dataset of quiet x-ray flares (i.e. 5843 x-ray flares with no proton associations).

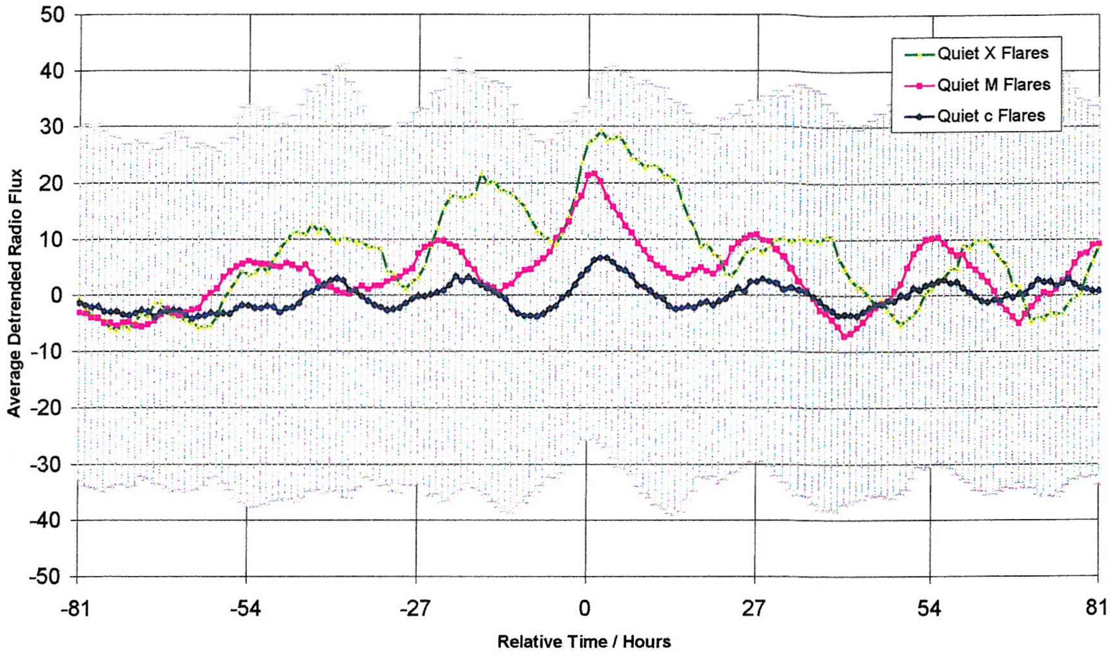

Figure 10-3 Variation in average detrended solar radio flux for quiet flares of class X, M and C. Error bars denote standard error and for clarity are shown only for the C-class case.

It is interesting to note that there is an apparent difference in the observed level of radio flux between flare classes, with X class flares correlating to higher levels of flux than M class flares, and M class flares correlating to higher levels of flux than C class flares, all of the order of 20 days prior to flare occurrence. The standard deviation is very large in proportion to the observed periodic variation, indicating that the use of the periodic and systematic variation as a predictor would be difficult, but the distributions do nevertheless show potential for a flare forecasting capability, relating higher radio flux levels to x-ray flares of greater magnitude.

10.6 Potential as a flare forecast model

The fact that the classification model was apparently triggered by x-ray flare 'precursors' suggested that it may have some use as a flare forecasting model as opposed to an SPE forecasting model. In order to test if this was true, a second analysis of the real-time POPE model forecast was performed by comparing the output of the model with the occurrence of x-ray flares.

10.6.1 Method

Over the period of real-time operation (December 2001 to December 2002) the GOES-8 hourly soft x-ray flux (i.e. the 1-8Å channel) was scanned and any hourly periods with a flux greater than $1\text{x}10^{-5}$ watts/m^2 were marked as 'x-ray events'. This flux threshold is equivalent to class M flares or greater. Hourly periods with a flux less than the threshold were marked as 'quiet'.

114 x-ray flare events were identified over this period, 93 of which occurred whilst POPE was running (for 21 of the x-ray flare events the POPE model was inactive, primarily due to drop-outs in internet access or intolerable gaps in GOES satellite data). Predictions from POPE were only considered if they were based on input delay vectors with no more than 50% missing data in each 3-hour interval. Hourly periods for which no GOES x-ray data was available were removed from the dataset. The prediction for each valid time interval was taken as the highest output from the model from within ± 2 hours of the current time.

Output from POPE consisted of an integer between 0 and 10, representing the number of networks from 10 that were predicting an event (in this case taken to be an x-ray flare of class M or above) in 48-hours time. A threshold analysis was performed in order to determine the optimal decision boundary, and skill scores were calculated from a contingency table. These have been compared against the POPE scores for SPE forecasting.

10.6.2 Results

The performance of the POPE multiple model is shown below in Table 10-9 for different threshold values.

Threshold	>5for flare	>6 for flare	>7 for flare	>8 for flare	>9 for flare
Flares from 93	78	75	70	69	39
QPs from 6699	2235	2839	3549	4385	5450
Flares / %	83.9	80.6	75.3	74.2	41.9
QPs / %	33.4	42.4	53.0	65.5	81.4
Average / %	**58.6**	**61.5**	**64.1**	**69.8**	**61.6**

Table 10-9 Classification success for real-time POPE model when asked to classify the occurrence of an M-class or greater x-ray event as opposed to an SPE.

As with the SPE detection model, optimal success occurs with a threshold of >8, meaning that 9 or 10 of the 10 constituent networks need to forecast an event (x-ray flare) before the overall POPE output is taken as being a positive-occurrence prediction. It can already be noted that the model has a 69.8% overall success rate which is greater than the 60.1% success rate of the SPE forecast model

Table 10-10 shows the contingency table for the >8 threshold and Table 10-11 shows the skill score for the POPE flare model in comparison to the POPE SPE model.

		Event Observed	
		Yes	No
Event Forecast	Yes	69 (74%)	2314 (35%)
	No	24 (26%)	4385 (65%)

Table 10-10 Contingency table for real-time operation of POPE model, Dec2001-Dec2002 for the forecast of M-class or greater x-ray flare events.

Statistic		POPE SPE	POPE Flare
Number of Hits ('Yes' forecast and 'Yes' observed)	A	6	69
False Alarms ('Yes' forecast but not observed).	B	2129	2314
Misses ('No' Forecast and 'Yes' observed)	C	5	24
Correct Nulls ('No' Forecast and 'No' observed)	D	4120	4385
Prob. of Detection 'Yes'	PODy	0.545	0.742
Prob. of Detection 'No'	PODn	0.659	0.655
False Alarm Ratio	FAR	0.997	0.971
Bias		194.091	25.624
Critical success index	CSI	0.003	0.029
True Skill Statistic	TSS	0.205	0.397
Heidke skill score	HSS	0.002	0.030
Gilbert skill score	GSS	0.001	0.015
Total number of events	N	6260	6792
Number expected correct by chance	C_1	4121.503	4381.259
Number of hits expected by chance	C_2	3.752	32.629

Table 10-11 Comparison of performance and skill statistics for POPE SPE forecast model and POPE as a >M class flare prediction model.

Examination of the skill scores shows that when operating as a flare prediction model the Heidke and Gilbert skill scores have a 10 fold improvement, hence the model does perform better, but the skill scores are still very low. The problem is again a relatively large proportion of false alarms.

Although the model correctly identifies 65.5% of quiet periods, the 34.5% that are miss-classified represent 2314 instances, which is very large compared to the actual number of observed flare-events (93), and it is this which ultimately dictates the skill scores. However, despite low skill scores the model is the only example of a 2-day autonomous flare forecast tool and may be useful for applications which have no serious consequences for false alarms.

10.7 Summary

The fact that SPEs have a strong correlation with solar x-ray flares means that neural models may have been trained to distinguish between cases of x-ray flare and no x-ray flare as opposed to SPE and no-SPE. This hypothesis was put forward as a way of explaining the high false alarm rate in the real-time POPE model, which forecast many more SPEs than were observed.

By considering x-ray flares that were observed by the SEC monitoring network between 1986 and 1999 it was shown that 92.9% of SPEs had flare associations, but only 0.7% of all x-ray flares were associated with SPEs. 32.4% of X class x-ray flares were found to be associated with SPEs even though this flare category only accounted for 0.8% of all flares. This indicates that if X class flares can be predicted accurately the probability of an SPE occurring is known. The same is also true for M class flares, although only 3.4% of M-class flares were found to have proton associations.

When the POPE model was queried with a dataset of over 5000 quiet periods selected so as to have an x-ray flare at time zero only 62% of the quiet periods were classified correctly. This compared to a 91% classification success when quiet periods were selected at random, indicating that it is highly likely that x-ray flares are causing the model to produce SPE forecasts.

Retraining the POPE networks using quiet period examples selected at the times of solar x-ray flares produced a 51% overall success rate showing that the neural models were unable to learn to differentiate between x-ray flares associated with SPEs and x-ray flares associated with quiet periods.

Reproducing the observed quiet period to SPE ratio in the training set did not produce an intelligent model, and simply caused quiet periods to be predicted for all input vectors.

An examination of the XS/XL ratio prior to quiet x-ray flares showed the average value to be somewhere between the cases of 'SPE' and quiet period. A similar examination of the 2800Mhz solar radio flux showed that prior to x-ray flares flux levels had a similar magnitude and periodicity to the SPE case. All evidence suggests that x-ray flares associated with quiet periods have a similar precursive behaviour in x-rays and radio flux to SPE associated flares.

Theorising that x-ray flares were the primary source of false alarms the POPE real-time model was reassessed as a forecaster of >M class x-ray flares. Overall classification success was measured at 69.8% and was greater than when operating as an SPE forecast model. This represents the only autonomous 2-day lead time flare forecast model in operation. Performance as a flare forecaster was still hampered by a high false alarm rate.

11. SUMMARY AND CONCLUSIONS

This section summarises the evolution of work contained in the thesis and highlights the primary findings. Avenues of future work are suggested and overall conclusions stated.

11.1 Summary

Current models to predict SPEs use characteristics of a discrete x-ray flare in order to produce a forecast, but the lead-times of such forecasts are limited in that significant proton flux can be observed within minutes of an x-ray flare peak. The need for forecasts with longer lead times may be required for manned missions of an interplanetary nature, and would serve as a useful tool for spacecraft operations in general.

This thesis addresses the fact that very little research has been conducted into developing SPE forecasting techniques that do not require an x-ray flare as a precursor. As an output, it has aimed to produce a forecast model with lead-times greater than that offered by current SPE prediction models.

Key to the development of a prediction methodology has been the use of a large dataset encompassing over 200 Solar Proton Events over a 3-decade time span. In addition, use has been made of GOES satellite observations to provide a uniform solar x-ray data source over a 25-year period, and ground station observations have provided continuous measurement of solar radio flux for over 50 years.

The work has succeeded in creating a real-time SPE forecast model with a 48-hour lead time that can operate autonomously using inputs from live GOES satellite measurements of solar x-ray flux. The technique uses a classification approach to the forecast problem and is based around the fact that solar x-ray fluxes are, on average, higher several days prior to SPEs than at times when no SPE occurs.

11.2 Main Findings

Application of an ARIMA time series forecasting method to the >10MeV daily proton flux time series showed that SPEs are not deterministic in the sense that their occurrence cannot be predicted from prior values of proton flux. This would appear to be the first time that this result has been formally documented.

Due to the discrete nature of solar proton events the solar proton flux time series can be approximated as a binary series composed of quiet periods and SPEs. In this way it has been possible to forecast SPEs using a classification approach.

Based on a sample of 97 SPEs the average GOES x-ray fluxes in both the XS(0.4-4Å) and XL(1-8Å) channels were found to rise from around 150 hours prior to the occurrence of an SPE. No rise in flux was observed in the cases of quiet periods. A comparison of the distributions showed them to be different in the mean to a 99% confidence level up to 150 hours prior to SPE occurrence. A similar finding was observed in the 2800Mhz solar radio flux.

The average solar radio flux associated with SPEs was found to contain a strong 27-day periodicity over a 162-day span centred on the event, with the occurrence of an SPE coinciding with a peak in the radio flux average. By contrast, quiet periods approximately coincided with a minimum in the average solar radio flux. Previous work by Chakravorti et al. [*Chakravorti, Das, Sen and Dasgupta, 1991*] has identified solar radio flux from active regions as passing through a maximum on the day of an SPE, but did not identify a 27-day periodicity in the flux. Evidence of a 27-day periodicity coinciding with SPE occurrence is also present to a lesser extent in the average solar x-ray flux. The correlation between SPEs and peaks in a 27-day periodicity has not been published elsewhere, but is almost certainly due to the fact that SPEs originate from active regions on the solar surface which evolve over a number of solar revolutions.

Knowing that solar radio flux from a source near the solar limb is significantly attenuated by the solar atmosphere, it is surprising that radio flux peaks are still seen at the time of an SPE; for an SPE to be monitored at earth it generally originates from a region towards the western solar limb. Use of spatially resolved radio flux measurements corrected for the observer location may therefore reveal a much higher correlation between solar radio flux peaks and SPE occurrence.

The development of a classification model proved that the differences between SPE and Quiet Period distributions in solar x-rays (specifically the solar x-ray ratio) could be harnessed to give a limited prediction capability with a 2-day lead time. During development, highest performance was given by an RBF model which classified 65% of examples correctly.

The output of the thesis has been POPE: a real-time variant of the classification model capable of autonomous operation on a rolling timeline with a 48-hour lead time. Assessing the model over a 12 month period showed it to have superior SPE detection capability to the Space Environment Centre's 2-day forecast model. POPE is the only autonomous SPE forecast model with a 48-hour lead time.

The SEC PROTONS model still represents the best in terms of SPE forecast accuracy. Whereas POPE predicted 55% of SPEs correctly during real-time operation, PROTONS is capable of forecasting 77% of SPEs successfully and has a much lower incidence of false alarms.

The operation of an SPE forecast model on a rolling timeline has been shown to be inherently problematic because significantly more quiet periods are observed than SPEs. As a result, the POPE model generates a large proportion of false alarms (during real time operation 99.7% of SPE forecasts were false alarms). In order to be effective in practice such a model needs to be able to identify quiet periods to a high degree of certainty.

The POPE model is superior to the SEC 2-Day model for applications in which it is more important to detect SPEs than to minimise false alarms. During real-time operation POPE can be expected to forecast 55% of SPEs and 66% of quiet periods successfully.

X-ray flares of class C and higher are also shown to be associated with increased levels of x-ray flux and solar radio flux several days prior to their occurrence. This is a new finding, but shows that increases in x-ray and radio flux are not specific to SPE occurrence. This hinders the prediction of SPEs when using the classification approach, as precursors to x-ray flares are mistakenly classified as SPEs by the model. This accounts for the high number of false alarms produced by the POPE classification method. It also means that precursors exist which potentially enable the forecasting of x-ray flares with lead-times of several days.

X class x-ray flares constitute only 0.8% of all flares yet 34% of them can be associated with SPEs. This means that the occurrence of x-ray flares can be used to infer the occurrence of an SPE. Specifically, if an X class x-ray flare can be forecast with certainty then there is a corresponding 34% probability of an SPE also occurring.

Analysis of the POPE model as a flare forecast tool showed superior performance to when acting as an SPE forecaster. During a 12-month period 74% of x-ray flares of class M0 or greater were forecast correctly, although a high false alarm rate was still observed in that 97% of flare predictions were false alarms.

11.3 Future Work

The limiting factor to the 2-day lead time classification technique has been the inseparability of the two classes in the predictor variables, whether this be solar x-ray fluxes, the x-ray ratio or solar radio fluxes. It stems from the fact that solar x-ray flares at the times of quiet periods exhibit similar precursors to the times of SPEs. If the technique is to be improved, it needs to incorporate knowledge that can allow more accurate classification of the predictor variables.

Combining forecaster experience with the model output may help to eliminate some false alarms when it is evident that there are no dangerous active regions on the solar disk, and this information could eventually be formalised as another input to the model.

One possible flaw with the current approach is that the inputs consist of flux-values and there is no feature extraction in terms of the dominant frequencies that are present in the solar x-ray or radio fluxes. Examination of the radio flux at the times of events revealed a strong near 27-day periodicity co-incident with the time of an SPE, which may be harnessable as a predictor, and to a lesser extent a similar pattern was seen in the x-ray flux. The obvious step is to use a Fourier analysis to identify frequencies within the input window, although the delay vectors would need to cover several solar rotations in order to reveal any 27-day periodicity. If SPEs are associated with active regions that last for about 4 solar rotations, this could be shown by correlating periods of 27-day periodicity with SPE occurrence via a wavelet analysis of the historical radio flux timeline. The problem with relying on periodicities as a predictor however is that they are a manifestation of a rotating active region, and if two or more active regions are present simultaneously, separated by varying degrees of longitude, it is likely to generate miscellaneous frequencies in the observed predictor variables. The spatial resolution of input data pertaining to specific active regions is therefore another desirable addition to the model, as this would allow separation of flux contributions from active regions that are present at the same time.

The current model utilises full disk solar measurements of x-ray flux, and although this is actually a good measure of active region activity, it cannot distinguish between different active regions. The ability to measure flux from specific active regions could yield a much higher correlation between predictor flux behaviour and SPE occurrence, and other spatial measurements, such as the heliolongitude of the active region, could be used as inputs. One problem that may be present in the current dataset is that in order to monitor an SPE at earth the particle injection site generally needs to be west of the solar central angle, whereas the electromagnetic nature of the predictor fluxes means that they are observed regardless of their source location. Some 'proton flares' may therefore have been observed as 'quiet periods' simply because the earth was not suitably magnetically connected to the particle injection site, but short of surrounding the sun with a ring of observing platforms it is difficult to see how this possibility might be addressed.

Whilst the full disk measurements utilised in this study cannot identify specific active regions, they have the advantage of being in a state of continuous monitoring, and the facility to have an uninterrupted observation should not be underestimated: especially if it is to be used as an input to a real-time model. Instruments providing spatially resolved measurements are unlikely to be dedicated and coverage of any one source will be intermittent. The lifetime of such instruments is also likely to be short compared to the observation span necessary to build a large dataset of SPE associated observations.

One of the most promising findings from the study was the fact that radio flux maxima coincided with the occurrence day of an SPE, despite the fact that the observed flux was probably significantly attenuated by the solar atmosphere. If spatially resolved radio flux measurements could be obtained for proton producing active regions, the observations could be corrected for heliographic position in relation to the solar central angle. The result may give a much stronger correlation between radio flux peaks and SPE occurrence with the potential to relate levels of solar radio flux quantitatively to SPE characteristics. Observations have already shown that more than 70% of SPEs occur within 6 days of a radio flux peak, and the nature of the radio flux variation gives the potential for forecasts with lead times of days.

11.4 Conclusions

SPEs are not deterministic when attempting to forecast them as part of a continuous proton flux time series, but can be approximated as binary events on a timeline. This lends itself well to forecasting via a classification approach.

On average, increases in solar x-ray flux, the solar x-ray ratio and the 2800Mhz solar radio flux are present from 150 hours prior to SPE occurrence. This allows the potential forecast of SPEs with a lead time of days. Increases in average x-ray and radio flux levels also occur prior to x-ray flares though and are not a unique precursor to SPEs.

A classification technique can correctly predict an SPE or quiet period 65% of the time with a 48-hour lead time using the $\log_{10}$ GOES XS/XL ratio as inputs over a 72-hour period.

The quiet period and SPE distributions in the x-ray and radio fluxes are significantly different in their means, but contain a high degree of scatter about their means that limits their seperability. The poor seperability of the predictor variables stems from the fact that x-ray flares at the times of quiet periods exhibit similar precursors to the times of SPEs, and it is this which limits the success of the classification technique.

A real-time classification model had superior SPE detection capability to the 2-day SPE forecast model operated by the Space Environment Centre and is the only autonomous real-time model in operation with a 48-hour lead time. It can be expected to forecast 55% of SPEs and 66% of Quiet Periods correctly.

As the POPE model requires no human supervision and little computational power, it is easy to install operationally and does not require large resources.

The concept of a rolling timeline SPE forecast necessitates the accurate prediction of Quiet Periods as these are seen far more frequently than SPEs. The real-time POPE model generates too many false alarms to be used as a reliable SPE forecaster, but is superior to the SEC 2-day forecast for applications in which false alarms can be tolerated.

Future work should try and correlate spatially resolved solar radio data to SPE occurrence. Correction of active region solar radio flux for centre-to-limb dependence may yield a higher correlation between radio flux maxima and SPEs, allowing the potential forecast of SPEs with several days lead time.

Appendix A
Listing of SPEs 1965 -1999

Appendix A-1

ID	SOURCE	Year	CYCLE	IMP Date	IMP Time	GOES Date	GOES Time	Log >10 MeV fluence	Log >30 MeV fluence	Log>60 MeV fluence	1st	End	Duration/ Days	GOES Class
1	IMP	1965	Min	06/02/65	Not found	NA	NA	7.216365	6.39658	5.877708	36	38	3	Not Found
2	IMP	1965	Min	06/10/65	Not found	NA	NA	6.110174	5.245238	4.698481	278	278	1	Not Found
3	IMP	1966	Min	25/03/66	Not found	NA	NA	6.979613	5.866953	5.166057	83	83	1	Not Found
4	IMP	1966	Min	08/07/66	Not found	NA	NA	7.799832	6.458149	5.688936	188	190	3	Not Found
5	IMP	1966	Min	28/08/66	Not found	NA	NA	7.700318	6.471249	5.725919	240	243	4	Not Found
6	IMP	1966	Min	03/09/66	Not found	NA	NA	8.982993	6.820341	5.873943	245	251	7	Not Found
7	IMP	1966	Min	15/09/66	Not found	NA	NA	7.220159	5.896061	5.096421	257	260	4	Not Found
8	IMP	1967	Max	13/01/67	Not found	NA	NA	6.411204	4.482882	3.336753	12	12	1	Not Found
9	IMP	1967	Max	29/01/67	Not found	NA	NA	9.041084	7.2031	6.33129	28	41	14	Not Found
10	IMP	1967	Max	15/02/67	Not found	NA	NA	6.316074	5.32128	4.698481	45	45	1	Not Found
11	IMP	1967	Max	01/03/67	Not found	NA	NA	6.615164	6.032235	5.666152	59	60	2	Not Found
12	IMP	1967	Max	13/03/67	Not found	NA	NA	7.147725	6.379722	5.899641	71	72	2	Not Found
13	IMP	1967	Max	25/05/67	Not found	NA	NA	8.888905	7.734024	7.159836	144	150	7	Not Found
14	IMP	1967	Max	07/06/67	Not found	NA	NA	7.191969	6.797125	6.605682	157	160	4	Not Found
15	IMP	1967	Max	15/06/67	Not found	NA	NA	6.447343	6.273518	6.255308	165	166	2	Not Found
16	IMP	1967	Max	03/11/67	Not found	NA	NA	6.863739	6.33041	6.218138	306	307	2	Not Found
17	IMP	1967	Max	13/11/67	Not found	NA	NA	7.145302	6.738068	6.697347	316	321	6	Not Found
18	IMP	1967	Max	04/12/67	Not found	NA	NA	7.378343	6.93538	6.654712	337	339	3	Not Found
19	IMP	1967	Max	17/12/67	Not found	NA	NA	7.104651	6.739874	6.585585	350	353	4	Not Found
20	IMP	1968	Max	18/07/68	Not found	NA	NA	6.092628	5.975243	5.9328	48	48	1	Not Found
21	IMP	1968	Max	28/03/68	Not found	NA	NA	7.389832	6.588391	6.589607	87	87	1	Not Found
22	IMP	1968	Max	27/04/68	Not found	NA	NA	6.501106	6.270747	6.227454	117	118	2	Not Found
23	IMP	1968	Max	10/06/68	Not found	NA	NA	8.46067	7.082219	6.691342	161	163	3	Not Found
24	IMP	1968	Max	08/07/68	Not found	NA	NA	7.658662	6.927539	6.839249	189	196	8	Not Found
25	IMP	1968	Max	27/09/68	Not found	NA	NA	7.622423	7.160521	6.947839	270	276	7	Not Found
26	IMP	1968	Max	05/10/68	Not found	NA	NA	7.497821	6.68855	6.369171	278	280	3	Not Found
27	IMP	1968	Max	30/10/68	Not found	NA	NA	6.23438	5.772916	5.750053	303	303	1	Not Found
28	IMP	1968	Max	01/11/68	Not found	NA	NA	8.312575	7.283353	6.712783	305	310	6	Not Found
29	IMP	1968	Max	19/11/68	Not found	NA	NA	9.006618	8.325674	7.560716	323	328	6	Not Found
30	IMP	1968	Max	04/12/68	Not found	NA	NA	8.361691	7.599726	6.984967	338	345	8	Not Found
31	IMP	1969	Max	24/01/69	Not found	NA	NA	6.26872	5.889421	5.848637	24	24	1	Not Found
32	IMP	1969	Max	26/02/69	Not found	NA	NA	7.873438	7.458379	7.1638	56	60	5	Not Found
33	IMP	1969	Max	22/03/69	Not found	NA	NA	6.794635	6.276272	6.147658	80	81	2	Not Found
34	IMP	1969	Max	31/03/69	Not found	NA	NA	9.362601	8.396566	7.750698	89	113	25	Not Found
35	IMP	1969	Max	26/09/69	Not found	NA	NA	7.251361	6.552391	6.466282	268	272	5	Not Found
36	IMP	1969	Max	15/10/69	Not found	NA	NA	6.169262	5.890636	5.78546	287	287	1	Not Found

37	IMP	1969	Max	03/11/69	Not found	NA	NA	8.804021	8.316406	7.677445	306	312	7	Not Found
38	IMP	1969	Max	25/11/69	Not found	NA	NA	6.531268	6.312185	6.134021	328	329	2	Not Found
39	IMP	1969	Max	21/12/69	Not found	NA	NA	6.384028	5.994287	5.80951	354	354	1	Not Found
40	IMP	1970	Max	29/01/70	Not found	NA	NA	7.431748	6.933955	6.64232	29	33	5	Not Found
41	IMP	1970	Max	07/03/70	Not found	NA	NA	7.82375	6.514578	6.357529	65	68	4	Not Found
42	IMP	1970	Max	14/03/70	Not found	NA	NA	6.714242	6.282714	6.114542	82	83	2	Not Found
43	IMP	1970	Max	17/03/70	Not found	NA	NA	6.479768	6.1385	6.081437	85	86	2	Not Found
44	IMP	1970	Max	20/03/70	Not found	NA	NA	7.903721	7.479033	7.095135	88	93	6	Not Found
45	IMP	1970	Max	31/05/70	Not found	NA	NA	7.082998	6.215275	6.078299	150	151	2	Not Found
46	IMP	1970	Max	27/06/70	Not found	NA	NA	6.468693	5.841224	5.74841	177	177	1	Not Found
47	IMP	1970	Max	08/07/70	Not found	NA	NA	6.538151	6.22437	6.066323	188	189	2	Not Found
48	IMP	1970	Max	24/07/70	Not found	NA	NA	7.53362	6.345836	6.203631	204	206	3	Not Found
49	IMP	1970	Max	14/08/70	Not found	NA	NA	8.268135	7.009128	6.734259	225	233	9	Not Found
50	IMP	1970	Max	06/11/70	Not found	NA	NA	7.806134	6.764726	6.549872	309	313	5	Not Found
51	IMP	1970	Max	14/12/70	Not found	NA	NA	6.445657	6.018092	6.097553	347	348	2	Not Found
52	IMP	1970	Max	25/12/70	Not found	NA	NA	7.134021	6.816976	6.651252	358	363	6	Not Found
53	IMP	1971	Max	25/01/71	Not found	NA	NA	9.171914	8.532643	7.800945	25	30	6	Not Found
54	IMP	1971	Max	03/03/71	Not found	NA	NA	6.20891	5.927261	5.8605	92	92	1	Not Found
55	IMP	1971	Max	07/03/71	Not found	NA	NA	7.49421	6.713512	6.414848	96	98	3	Not Found
56	IMP	1971	Max	22/04/71	Not found	NA	NA	6.223244	5.927818	5.858545	111	111	1	Not Found
57	IMP	1971	Max	17/05/71	Not found	NA	NA	6.958449	6.364714	6.187093	136	137	2	Not Found
58	IMP	1971	Max	02/09/71	Not found	NA	NA	8.581896	8.2233	7.803295	244	251	8	Not Found
59	IMP	1971	Max	05/10/71	Not found	NA	NA	6.701304	6.518597	6.316757	277	278	2	Not Found
60	IMP	1972	Max	07/03/72	Not found	NA	NA	6.410471	6.069147	6.010695	66	66	1	Not Found
61	IMP	1972	Max	19/04/72	Not found	NA	NA	7.450195	6.614363	6.510376	109	111	3	Not Found
62	IMP	1972	Max	29/04/72	Not found	NA	NA	7.874636	7.094529	6.921141	149	155	7	Not Found
63	IMP	1972	Max	09/06/72	Not found	NA	NA	6.977731	6.618922	6.497222	160	162	3	Not Found
64	IMP	1972	Max	17/06/72	Not found	NA	NA	7.362469	6.566051	6.468212	168	170	3	Not Found
65	IMP	1972	Max	20/07/72	Not found	NA	NA	10.05404	9.697925	8.351868	201	234	34	Not Found
66	IMP	1972	Max	31/10/72	Not found	NA	NA	7.540873	6.518597	6.367959	304	305	2	Not Found
67	IMP	1972	Max	11/12/72	Not found	NA	NA	7.017089	5.540873	5.536783	345	345	1	Not Found
68	IMP	1973	Max	13/04/73	Not found	NA	NA	6.750053	6.08957	5.943134	102	103	2	Not Found
69	IMP	1973	Max	01/05/73	Not found	NA	NA	7.052339	6.86419	6.72236	120	122	3	Not Found
70	IMP	1973	Max	31/07/73	Not found	NA	NA	6.725032	6.395749	6.266683	211	213	3	Not Found
71	IMP	1973	Max	08/09/73	Not found	NA	NA	7.313103	6.63221	6.453195	250	253	4	Not Found
72	IMP	1973	Max	04/11/73	Not found	NA	NA	6.437124	5.970222	5.755054	307	307	1	Not Found
73	IMP	1974	Min	01/01/74	Not found	NA	NA	7.65062	7.641028	7.603925	1	1	1	Not Found
74	IMP	1974	Min	15/01/74	Not found	NA	NA	7.084941	7.084941	7.084941	15	15	1	Not Found

75	IMP	1974	Min	10/06/74	Not found	NA	NA	6.384028	5.572282	5.565923	160	160	1	Not Found
76	IMP	1974	Min	03/07/74	6	03/07/74	18	8.519125	7.557476	7.409665	184	189	6	1
77	IMP	1974	Min	11/09/74	0	11/09/74	2	8.308539	7.458706	7.185773	254	260	7	1
78	IMP	1974	Min	20/09/74	0	20/09/74	0	8.090144	7.322897	7.152862	263	271	9	1
79	IMP	1974	Min	05/11/74	14	05/11/74	16	7.084941	6.520023	6.276522	309	310	2	1
80	IMP	1975	Min	21/08/75	14	no data	no data	6.739015	6.344927	6.17371	233	234	2	5
81	IMP	1976	Min	30/04/76	18	30/04/76	22	8.005139	7.545866	7.286996	121	123	3	1
82	IMP	1976	Min	22/08/76	12	22/08/76	14	6.951651	6.377158	6.174973	235	236	2	1
83	IMP	1977	Min	09/09/77	no data	09/09/77	8	7.217567	6.567713	6.48797	252	257	6	5
84	IMP	1977	Min	17/09/77	0	19/09/77	0	8.470643	7.672291	7.407253	260	264	5	1
85	IMP	1977	Min	24/09/77	6	24/09/77	8	7.924745	7.481592	7.274297	267	270	4	1
86	IMP	1977	Min	12/10/77	2	12/10/77	4	6.406791	5.96769	5.825304	285	285	1	3
87	IMP	1977	Max	22/11/77	10	22/11/77	10	8.47042	7.913751	7.659076	326	329	4	1
88	IMP	1978	Max	02/01/78	0	no data	no data	6.833683	6.431748	6.293642	2	4	3	5
89	IMP	1978	Max	13/02/78	6	13/02/78	8	9.168236	8.263391	7.982632	44	49	6	1
90	IMP	1978	Max	08/04/78	4	08/04/78	4	6.114905	5.671207	5.610911	98	98	1	5
91	IMP	1978	Max	11/04/78	14	11/04/78	16	7.80377	7.205398	6.891545	101	103	3	1
92	IMP	1978	Max	17/04/78	2	17/04/78	8	9.383636	8.568271	8.267117	107	129	23	5
93	IMP	1978	Max	31/05/78	no data	31/05/78	12	7.221549	6.260774	6.1834	151	153	3	1
94	IMP	1978	Max	23/06/78	0	23/06/78	12	7.711227	6.744399	6.657626	174	177	4	2
95	IMP	1978	Max	12/07/78	12	13/07/78	0	7.343005	6.532238	6.414121	193	197	5	1
96	IMP	1978	Max	08/09/78	2	08/09/78	2	6.261033	5.735561	5.63232	251	251	1	5
97	IMP	1978	Max	23/09/78	10	23/09/78	12	9.451284	8.708781	8.338565	266	271	6	1
98	IMP	1978	Max	08/10/78	20	09/10/78	22	6.742441	6.144289	5.987061	282	283	2	1
99	IMP	1978	Max	10/11/78	2	10/11/78	6	7.218709	6.269987	6.226895	314	316	3	2
100	IMP	1978	Max	11/12/78	22	12/12/78	0	6.703176	6.070753	6.050663	346	348	3	5
101	IMP	1979	Max	17/02/79	20	17/02/79	20	7.150001	6.597301	6.433316	48	51	4	1
102	IMP	1979	Max	02/03/79	0	02/03/79	12	7.060209	6.526244	6.429649	61	65	5	5
103	IMP	1979	Max	10/03/79	12	no event	no event	6.37418	5.898451	5.847968	70	71	2	5
104	IMP	1979	Max	14/03/79	18	14/03/79	20	6.156297	5.567203	5.548941	74	74	1	5
105	IMP	1979	Max	03/04/79	6	03/04/79	16	7.311726	6.300541	6.217853	93	95	3	1
106	IMP	1979	Max	06/06/79	16	06/06/79	18	8.42835	7.338811	7.254075	157	160	4	1
107	IMP	1979	Max	05/07/79	22	05/07/79	20	7.282714	6.29651	6.200076	187	189	3	3
108	IMP	1979	Max	01/08/79	12	no event	no event	6.52266	6.32218	6.203041	218	220	3	5
109	IMP	1979	Max	19/08/79	10	19/08/79	8	8.789176	8.059016	7.785923	231	241	11	1
110	IMP	1979	Max	08/09/79	10	08/09/79	12	6.501106	6.082609	5.919952	251	252	2	3
111	IMP	1979	Max	10/09/79	20	no event	no event	6.471886	6.038753	5.891243	254	255	2	5
112	IMP	1979	Max	14/09/79	12	15/09/79	10	8.545221	8.113236	7.948849	258	271	14	2

113	IMP	1979	Max	16/11/79	0	16/11/79	0	7.516448	6.539378	6.374379	320	321	2	1
114	IMP	1980	Max	11/01/80	6	no event	no event	6.223244	5.496621	5.48443	11	11	1	5
115	IMP	1980	Max	06/02/80	no data	06/02/80	12	6.169262	5.610911	5.489042	38	38	1	1
116	IMP	1980	Max	04/04/80	16	04/04/80	16	6.716965	6.151999	5.986088	96	97	2	1
117	IMP	1980	Max	17/07/80	22	18/07/80	0	8.054215	7.12897	7.053133	200	203	4	1
118	IMP	1980	Max	23/07/80	6	continuation	continuation	6.771322	5.974243	5.935544	205	207	3	5
119	IMP	1980	Max	06/08/80	18	06/08/80	22	6.295794	5.705969	5.529878	220	220	1	5
120	IMP	1980	Max	15/10/80	10	15/10/80	12	7.446332	6.690093	6.520308	289	294	6	1
121	IMP	1980	Max	14/11/80	22	14/11/80	22	6.643178	6.11054	5.958449	320	322	3	5
122	IMP	1980	Max	23/11/80	20	23/11/80	22	6.828815	6.040474	5.861798	329	330	2	1
123	IMP	1980	Max	30/11/80	0	no event	no event	6.295794	5.462234	5.42133	335	335	1	5
124	IMP	1980	Max	02/12/80	0	no event	no event	6.077116	5.357943	5.341075	337	337	1	5
125	IMP	1981	Max	18/02/81	no event	no data	no data	8.823184	8.823184	8.823184	49	49	1	5
126	IMP	1981	Max	30/03/81	8	30/03/81	8	7.434704	6.786771	6.611258	89	95	7	1
127	IMP	1981	Max	10/04/81	18	10/04/81	18	7.83918	7.246844	6.975243	100	102	3	1
128	IMP	1981	Max	15/04/81	18	16/04/81	0	7.073548	6.26361	6.184635	105	108	4	4
129	IMP	1981	Max	24/04/81	14	no data	no data	9.039925	8.359314	8.2215	114	147	34	5
130	IMP	1981	Max	20/07/81	14	20/07/81	14	7.857957	7.112072	6.787541	201	203	3	1
131	IMP	1981	Max	24/07/81	16	24/07/81	18	6.876457	5.921649	5.86567	205	206	2	4
132	IMP	1981	Max	10/08/81	0	09/08/81	18	7.066323	6.051083	5.952177	221	222	2	2
133	IMP	1981	Max	06/09/81	20	no data	no data	6.796146	5.821053	5.813875	250	251	2	5
134	IMP	1981	Max	19/09/81	0	19/09/81	4	6.986575	6.325089	6.174973	262	265	4	5
135	IMP	1981	Max	24/09/81	0	no data	no data	6.107605	5.48751	5.437124	267	267	1	5
136	IMP	1981	Max	08/10/81	2	08/10/81	2	9.313321	8.669065	8.370259	281	294	14	1
137	IMP	1981	Max	10/11/81	no data	10/11/81	18	6.341075	5.438844	5.403079	315	315	1	3
138	IMP	1981	Max	15/11/81	0	15/11/81	2	6.061029	5.518597	5.397451	319	319	1	5
139	IMP	1981	Max	23/11/81	no data	23/11/81	14	6.217567	5.47174	5.391749	328	328	1	3
140	IMP	1981	Max	05/12/81	18	05/12/81	20	6.702241	5.827414	5.781579	340	341	2	3
141	IMP	1981	Max	10/12/81	4	10/12/81	2	7.837812	6.781423	6.601099	344	345	2	1
142	IMP	1981	Max	29/12/81	no data	27/12/81	14	6.788539	5.847968	5.81315	362	363	2	1
143	IMP	1982	Max	02/01/82	no data	02/01/82	8	8.471886	8.470292	8.470292	2	2	1	5
144	IMP	1982	Max	31/01/82	no data	31/01/82	2	9.075352	8.292377	7.987595	31	39	9	1
145	IMP	1982	Max	07/03/82	2	07/03/82	4	7.096798	6.257399	6.115628	66	67	2	1
146	IMP	1982	Max	30/03/82	no data	31/03/82	4	6.084941	5.435397	5.410472	90	90	1	5
147	IMP	1982	Max	04/06/82	12	07/06/82	0	7.815969	7.326337	7.184141	156	168	13	2
148	IMP	1982	Max	18/06/82	20	continuation	continuation	6.044323	5.499616	5.454025	170	170	1	5
149	IMP	1982	Max	28/06/82	no data	28/06/82	4	6.5748	6.200076	6.000925	179	179	1	3
150	IMP	1982	Max	10/07/82	0	11/07/82	2	8.845667	7.881014	7.703745	191	200	10	2

151	IMP	1982	Max	22/07/82	22	22/07/82	20	8.053132	7.101266	6.852428	204	207	4	1
152	IMP	1982	Max	04/09/82	12	05/09/82	0	7.133674	6.034419	5.987546	248	249	2	2
153	IMP	1982	Max	21/10/82	no event	no data	no data	6.856581	6.855268	6.855268	294	294	1	5
154	IMP	1982	Max	24/10/82	bad data	no data	no data	6.330189	6.309881	6.309881	297	297	1	5
155	IMP	1982	Max	22/11/82	18	22/11/82	20	8.999072	8.292955	8.009654	326	355	30	1
156	IMP	1983	Max	25/12/82	12	26/12/82	6	8.332892	7.453991	7.236601	360	0	7	5
157	IMP	1983	Max	03/02/83	8	03/02/83	8	8.013904	6.897616	6.789536	34	36	3	2
158	IMP	1983	Max	15/04/83	10	15/04/83	12	6.139527	5.891243	5.680162	105	105	1	5
159	IMP	1983	Max	15/06/83	0	15/06/83	4	7.133674	6.627567	6.369171	166	168	3	1
160	IMP	1984	Max	31/01/84	no data	31/01/84	16	6.092628	5.489042	5.458969	32	32	1	1
161	IMP	1984	Max	16/02/84	10	16/02/84	8	8.13826	7.404549	7.504221	47	55	9	1
162	IMP	1984	Max	12/03/84	0	12/03/84	0	7.473633	6.828885	6.606383	72	77	6	2
163	IMP	1984	Max	25/04/84	no data	25/04/84	8	9.118032	8.571347	8.444193	116	127	12	1
164	IMP	1984	Max	13/05/84	0	12/05/84	22	6.253207	5.485972	5.479768	134	135	2	5
165	IMP	1984	Max	24/05/84	0	24/05/84	6	6.666151	5.805839	5.78854	145	146	2	1
166	IMP	1985	Min	22/01/85	0	22/01/85	2	6.897258	6.368566	6.123859	22	23	2	1
167	IMP	1985	Min	24/04/85	12	24/04/85	10	7.908637	6.990207	6.834719	114	117	4	1
168	IMP	1985	Min	09/07/85	0	09/07/85	2	7.27377	6.624555	6.343219	190	190	1	1
169	IMP	1986	Min	04/02/86	8	04/02/86	10	8.130089	7.597659	7.437589	36	42	7	2
170	GOES	1986	Min	NA	NA	06/02/86	8	not found	not found	not found	NA	NA	4	1
171	IMP	1986	Min	14/02/86	10	14/02/86	14	8.430577	7.63112	7.356098	45	49	5	1
172	IMP	1986	Min	06/03/86	16	06/03/86	18	6.596825	6.011155	5.930039	65	66	2	1
173	IMP	1986	Min	04/05/86	10	04/05/86	12	6.205985	5.755882	5.643178	125	125	1	1
174	IMP	1987	Min	07/11/87	22	08/11/87	0	7.381685	6.432271	6.303833	312	313	2	1
175	IMP+GOES	1988	Max	30/12/87	2	29/12/87	22	8.000643	7.073787	6.824245	364	6	9	1
176	GOES	1988	Max	NA	NA	02/01/88	23	not found	not found	not found	NA	NA	6	1
177	IMP+GOES	1988	Max	25/03/88	22	25/03/88	23	6.726805	6.216709	6.046447	85	86	2	1
178	IMP	1988	Max	26/03/88	no event	no event	no event	6.788539	5.518597	5.51429	109	109	1	5
179	IMP + GOES	1988	Max	30/06/88	4	30/06/88	10	6.594432	5.995718	5.843935	182	182	1	1
180	IMP + GOES	1988	Max	24/08/88	0	24/08/88	5	7.154319	6.401211	6.368364	238	243	6	2
181	IMP+GOES	1988	Max	05/10/88	0	04/10/88	22	6.239843	5.563354	5.534034	279	279	1	3
182	IMP+GOES	1988	Max	12/10/88	4	12/10/88	6	6.146313	5.730329	5.608595	288	288	1	1
183	IMP+GOES	1988	Max	08/11/88	12	08/11/88	14	6.847968	6.403265	6.195591	314	315	2	1
184	GOES	1988	Max	NA	NA	14/11/88	0	not found	not found	not found	NA	NA	1	3
185	IMP+GOES	1988	Max	14/12/88	14	14/12/88	13	7.84797	7.086645	6.951492	350	356	7	1
186	IMP	1988	Max	27/12/88	no event	27/12/88	no event	6.081046	5.723252	5.657938	361	361	1	1
187	IMP+GOES	1989	Max	04/01/89	20	04/01/89	23	6.415935	5.641028	5.588392	5	5	1	1
188	GOES	1989	max	NA	NA	18/01/89	9	not found	not found	not found	NA	NA	1	3

189	IMP	1989	Max	03/02/89	no event?	data gap	data gap	7.891243	5.442264	5.431923	34	34	1	5
190	IMP+GOES	1989	Max	07/03/89	0	07/03/89	17	8.577053	7.621657	7.43056	67	75	9	2
191	IMP+GOES	1989	Max	17/03/89	14	17/03/89	19	8.850903	7.846762	7.661324	77	79	3	1
192	IMP+GOES	1989	Max	23/03/89	20	23/03/89	20	7.276272	6.451531	6.239572	82	84	3	1
193	IMP+GOES	1989	Max	09/04/89	18	10/04/89	18	8.213692	7.230321	7.102236	101	106	6	2
194	IMP+GOES	1989	Max	22/04/89	2	22/04/89	21	6.696589	6.386359	6.151334	113	115	3	2
195	GOES	1989	Max	NA	NA	01/05/89	19	not found	not found	not found	NA	NA	9	3
196	IMP+GOES	1989	Max	05/05/89	continuation	04/05/89	continuation	7.513856	6.645424	6.564512	121	128	8	4
197	IMP+GOES	1989	Max	20/05/89	14	20/05/89	13	7.395938	6.801764	6.699707	141	148	8	1
198	IMP+GOES	1989	Max	18/06/89	14	18/06/89	15	6.587173	6.196492	5.986575	169	170	2	1
199	IMP+GOES	1989	Max	29/06/89	2	29/06/89	4	6.683106	5.946881	5.910205	180	182	3	1
200	GOES	1989	Max	NA	NA	29/06/89	23	not found	not found	not found	NA	NA	2	3
201	IMP+GOES	1989	Max	25/07/89	data gap	25/07/89	8	7.172127	6.718771	6.443624	206	207	2	1
202	IMP+GOES	1989	Max	12/08/89	data gap	12/08/89	15	9.838331	9.258839	8.920182	225	249	25	1
203	IMP+GOES	1989	Max	12/09/89	10	12/09/89	14	7.434877	6.657419	6.452696	255	259	5	1
204	IMP	1989	Max	18/09/89	continuation	18/09/89	0	6.129145	5.560768	5.449023	261	261	1	1
205	IMP+GOES	1989	Max	29/09/89	12	29/09/89	12	9.532503	9.111428	8.848455	272	288	17	1
206	IMP+GOES	1989	Max	19/10/89	12	19/10/89	13	10.1186	9.675262	9.471343	292	313	22	1
207	IMP+GOES	1989	Max	15/11/89	6	15/11/89	7	7.484893	7.050663	6.814019	319	324	6	1
208	IMP+GOES	1989	Max	18/11/89	18	18/11/89	19	6.069147	5.613215	5.499616	326	326	1	1
209	IMP+GOES	1989	Max	26/11/89	12	26/11/89	21	9.141598	8.306908	8.027667	331	339	9	2
210	GOES	1989	Max	NA	NA	30/11/89	14	Not found	Not found	Not found	NA	NA	6	1
211	GOES	1990	Max	NA	NA	19/03/90	6	Not found	Not found	Not found	NA	NA	3	1
212	GOES	1990	Max	NA	NA	28/03/90	15	Not found	Not found	Not found	NA	NA	3	1
213	GOES	1990	Max	NA	NA	07/04/90	7	Not found	Not found	Not found	NA	NA	3	1
214	GOES	1990	Max	NA	NA	16/04/90	0	Not found	Not found	Not found	NA	NA	8	1
215	GOES	1990	Max	NA	NA	28/04/90	5	Not found	Not found	Not found	NA	NA	2	1
216	GOES	1990	Max	NA	NA	07/05/90	10	Not found	Not found	Not found	NA	NA	6	3
217	GOES	1990	Max	NA	NA	15/05/90	18	Not found	Not found	Not found	NA	NA	7	3
218	GOES	1990	Max	NA	NA	21/05/90	23	Not found	Not found	Not found	NA	NA	3	1
219	GOES	1990	Max	NA	NA	24/05/90	21	Not found	Not found	Not found	NA	NA	9	1
220	GOES	1990	Max	NA	NA	12/06/90	8	Not found	Not found	Not found	NA	NA	3	1
221	GOES	1990	Max	NA	NA	13/07/90	18	Not found	Not found	Not found	NA	NA	2	1
222	GOES	1990	Max	NA	NA	26/07/90	1	Not found	Not found	Not found	NA	NA	5	1
223	GOES	1990	Max	NA	NA	31/07/90	0	Not found	Not found	Not found	NA	NA	7	2
224	GOES	1990	Max	NA	NA	23/10/90	11	Not found	Not found	Not found	NA	NA	2	3
225	GOES	1990	Max	NA	NA	11/11/90	6	Not found	Not found	Not found	NA	NA	1	3
226	GOES	1990	Max	NA	NA	23/12/90	11	Not found	Not found	Not found	NA	NA	3	3

227	GOES	1991	Max	NA	NA	25/01/91	21	Not found	Not found	Not found	NA	NA	9	2
228	GOES	1991	Max	NA	NA	08/02/91	10	Not found	Not found	Not found	NA	NA	2	1
229	GOES	1991	Max	NA	NA	25/02/91	12	Not found	Not found	Not found	NA	NA	2	1
230	GOES	1991	Max	NA	NA	23/03/91	8	Not found	Not found	Not found	NA	NA	8	1
231	GOES	1991	Max	NA	NA	31/03/91	22	Not found	Not found	Not found	NA	NA	3	3
232	GOES	1991	Max	NA	NA	03/04/91	0	Not found	Not found	Not found	NA	NA	6	4
233	GOES	1991	Max	NA	NA	10/05/91	15	Not found	Not found	Not found	NA	NA	2	3
234	GOES	1991	Max	NA	NA	13/05/91	3	Not found	Not found	Not found	NA	NA	3	1
235	GOES	1991	Max	NA	NA	18/05/91	16	Not found	Not found	Not found	NA	NA	10	1
236	GOES	1991	Max	NA	NA	31/05/91	9	Not found	Not found	Not found	NA	NA	21	3
237	GOES	1991	Max	NA	NA	15/06/91	9	Not found	Not found	Not found	NA	NA	7	1
238	GOES	1991	Max	NA	NA	30/06/91	0	Not found	Not found	Not found	NA	NA	14	2
239	GOES	1991	Max	NA	NA	07/07/91	5	Not found	Not found	Not found	NA	NA	7	1
240	GOES	1991	Max	NA	NA	25/08/91	15	Not found	Not found	Not found	NA	NA	6	2
241	GOES	1991	Max	NA	NA	30/09/91	6	Not found	Not found	Not found	NA	NA	3	3
242	GOES	1991	Max	NA	NA	28/10/91	6	Not found	Not found	Not found	NA	NA	2	1
243	GOES	1991	Max	NA	NA	30/10/91	7	Not found	Not found	Not found	NA	NA	2	1
244	GOES	1992	Max	NA	NA	06/02/92	22	Not found	Not found	Not found	NA	NA	3	1
245	GOES	1992	Max	NA	NA	27/02/92	12	Not found	Not found	Not found	NA	NA	2	3
246	GOES	1992	Max	NA	NA	07/03/92	14	Not found	Not found	Not found	NA	NA	5	3
247	GOES	1992	Max	NA	NA	15/03/92	9	Not found	Not found	Not found	NA	NA	3	3
248	GOES	1992	Max	NA	NA	09/05/92	0	Not found	Not found	Not found	NA	NA	4	2
249	GOES	1992	Max	NA	NA	25/06/92	21	Not found	Not found	Not found	NA	NA	7	1
250	GOES	1992	Max	NA	NA	04/08/92	15	Not found	Not found	Not found	NA	NA	4	1
251	GOES	1992	Max	NA	NA	30/10/92	19	Not found	Not found	Not found	NA	NA	10	1
252	GOES	1992	Max	NA	NA	23/11/92	19	Not found	Not found	Not found	NA	NA	2	3
253	GOES	1992	Max	NA	NA	29/11/92	11	Not found	Not found	Not found	NA	NA	2	3
254	GOES	1993	Max	NA	NA	04/03/93	13	Not found	Not found	Not found	NA	NA	3	3
255	GOES	1993	Max	NA	NA	06/03/93	22	Not found	Not found	Not found	NA	NA	4	3
256	GOES	1993	Max	NA	NA	12/03/93	19	Not found	Not found	Not found	NA	NA	3	1
257	GOES	1993	Max	NA	NA	15/05/93	0	Not found	Not found	Not found	NA	NA	2	3
258	GOES	1993	Max	NA	NA	07/06/93	16	Not found	Not found	Not found	NA	NA	2	1
259	GOES	1993	Max	NA	NA	24/06/93	20	Not found	Not found	Not found	NA	NA	5	3+2
260	GOES	1993	Max	NA	NA	25/09/93	4	Not found	Not found	Not found	NA	NA	1	5
261	GOES	1994	Max	NA	NA	20/02/94	2	Not found	Not found	Not found	NA	NA	3	1
262	GOES	1994	Min	NA	NA	19/10/94	22	Not found	Not found	Not found	NA	NA	2	1
263	GOES	1995	Min	NA	NA	22/04/95	14	Not found	Not found	Not found	NA	NA	1	3
264	GOES	1995	Min	NA	NA	20/10/95	6	Not found	Not found	Not found	NA	NA	2	1

265	GOES	1997	Min	NA	NA	04/11/97	6	Not found	Not found	Not found	NA	NA	3	1
266	GOES	1997	Min	NA	NA	06/11/97	12	Not found	Not found	Not found	NA	NA	6	4
267	GOES	1997	Min	NA	NA	14/11/97	0	Not found	Not found	Not found	NA	NA	1	3
268	GOES	1998	Max	NA	NA	20/04/98	12	Not found	Not found	Not found	NA	NA	7	1
269	GOES	1998	Max	NA	NA	30/04/98	2	Not found	Not found	Not found	NA	NA	2	3
270	GOES	1998	Max	NA	NA	02/05/98	14	Not found	Not found	Not found	NA	NA	3	1
271	GOES	1998	Max	NA	NA	06/05/98	8	Not found	Not found	Not found	NA	NA	2	1
272	GOES	1998	Max	NA	NA	09/05/98	6	Not found	Not found	Not found	NA	NA	3	1
273	GOES	1998	Max	NA	NA	16/06/98	4	Not found	Not found	Not found	NA	NA	3	3
274	GOES	1998	Max	NA	NA	22/08/98	10	Not found	Not found	Not found	NA	NA	3	3
275	GOES	1998	Max	NA	NA	24/08/98	23	Not found	Not found	Not found	NA	NA	8	1
276	GOES	1998	Max	NA	NA	24/09/98	19	Not found	Not found	Not found	NA	NA	2	1
277	GOES	1998	Max	NA	NA	30/09/98	15	Not found	Not found	Not found	NA	NA	5	1
278	GOES	1998	Max	NA	NA	19/10/98	0	Not found	Not found	Not found	NA	NA	2	3
279	GOES	1998	Max	NA	NA	06/11/98	6	Not found	Not found	Not found	NA	NA	3	1
280	GOES	1998	Max	NA	NA	14/11/98	7	Not found	Not found	Not found	NA	NA	4	1
281	GOES	1999	Max	NA	NA	20/01/99	22	Not found	Not found	Not found	NA	NA	6	1
282	GOES	1999	Max	NA	NA	24/04/99	16	Not found	Not found	Not found	NA	NA	3	1
283	GOES	1999	Max	NA	NA	04/05/99	0	Not found	Not found	Not found	NA	NA	4	2
284	GOES	1999	Max	NA	NA	01/06/99	22	Not found	Not found	Not found	NA	NA	7	1

Appendix B
X-Ray Flare Class Associations for SPEs

ID	GOES Date	GOES Time	Duration / Days	Class	cycle	Start Time	Max Time	End Time	Long	Lat	Optical
1	29/12/87	22	9	1	Max	2322	2353	2322	S36	E34	SF
2	02/01/88	23	6	1	Max	2111	2409	2135	S34	W18	3B
3	25/03/88	23	2	1	Max	2141	2220	2154			
4	30/06/88	10	1	1	Max	904	916	906	S16	E22	2B
5	24/08/88	5	6	2	Max	1757E	1831D	1800	N24	E90	
6	04/10/88	22	1	3	Max	1226	1254	1238			
7	12/10/88	6	1	1	Max	0457E	0610D	502	S20	W66	2N
8	08/11/88	14	2	1	Max	1228	1414D	1248U	N16	W07	2F
9	14/11/88	0	1	3	Max	2252E	2405D	2309	S23	W27	1N
10	14/12/88	13	7	1	Max	1337E	1354D	1337	N30	E63	1B
11	04/01/89	23	1	1	Max	1738	1748D	1738U	N24	E67	SF
12	18/01/89	9	1	3	max	702	713	0707U	S30	W65	1F
13	07/03/89	17	9	2	Max	1447E	1545D	1454	N32	E65	2B
14	17/03/89	19	3	1	Max	1729E	1932D	1737	N33	W60	2B
15	23/03/89	20	3	1	Max	1925	2126	1948U	N18	W28	3B
16	10/04/89	18	6	2	Max	0044E	0239D	59	N35	E29	4B
17	22/04/89	21	3	2	Max	247	255	252			
18	01/05/89	19	9	3	Max	108	115	0110U	N28	E60	SN
19	20/05/89	13	8	1	Max	929	1308	1155			
20	18/06/89	15	2	1	Max	NO FLARE ASSOCIATION					
21	29/06/89	4	3	1	Max	0254E	0510D	315	N28	W28	1N
22	29/06/89	23	2	3	Max	2108	2350	2118U	N28	W44	2B
23	25/07/89	8	2	1	Max	0839E	0908D	843	N25	W84	2N
24	12/08/89	15	25	1	Max	1357E	1611D	1424	S16	W37	2B
25	12/09/89	14	5	1	Max	435	1057	814			
26	18/09/89	0	1	1	Max	1831E	1926D	1848	N14	W54	1F
27	29/09/89	12	17	1	Max	1047	1435	1133			
28	19/10/89	13	22	1	Max	1229E	2013	1255	S27	E10	4B
29	15/11/89	7	6	1	Max	0638E	0920D	705	N11	W26	3B
30	18/11/89	19	1	1	Max	1551E	1900D	1639	S30	W16	3B
31	26/11/89	21	9	2	Max	1756E	2240	1931	N25	W03	2B
32	30/11/89	14	6	1	Max	1145E	1504	1225	N24	W52	3B
33	19/03/90	6	3	1	Max	0439E	0620D	508	N31	W43	2B
34	28/03/90	15	3	1	Max	0731E	0948D	745	S04	W37	2N
35	07/04/90	7	3	1	Max	NO FLARE ASSOCIATION					
36	16/04/90	0	8	1	Max	255	345	0255U	N32	E57	2B
37	28/04/90	5	2	1	Max	2227	2438	2426			
38	07/05/90	10	6	3	Max	54	0104D	0055U	S06	W83	SF
39	15/05/90	18	7	3	Max	1246E	1613D	1310	N42	E38	3B
40	21/05/90	23	3	1	Max	2212E	2339D	2217	N35	W36	2B
41	24/05/90	21	9	1	Max	2046E	2145D	2049	N33	W78	1B
42	12/06/90	8	3	1	Max	455	0735D	529	N10	W33	2B
43	13/07/90	18	2	1	Max	959	1242	1130			
44	26/07/90	1	5	1	Max	2223E	2529D	2324	S15	E55	2N
45	31/07/90	0	7	2	Max	0633E	0847D	732	N20	E45	2B
46	23/10/90	11	2	3	Max	NO FLARE ASSOCIATION					
47	11/11/90	6	1	3	Max	0446E	0454D	449	N10	E18	SF
48	23/12/90	11	3	3	Max	943	1029D	1001	N11	W68	2B
49	25/01/91	21	9	2	Max	0630E	0638D	630	S16	E78	SF
50	08/02/91	10	2	1	Max	629	911	748			
51	25/02/91	12	2	1	Max	0809E	0930D	822	S16	W80	2N
52	23/03/91	8	8	1	Max	2243E	2317D	2245	S26	E28	3B

53	31/03/91	22	3	3	Max	1927	2359	2034			
54	10/05/91	15	2	3	Max	0254E	357	346	N39	E10	SF
55	13/05/91	3	3	1	Max	103	222	144			
56	18/05/91	16	10	1	Max	0506E	0748D	546	N32	W85	2B
57	31/05/91	9	21	3	Max	2214E	2443D	2343	N05	E38	2B
58	15/06/91	9	7	1	Max	0633E	1117D	831	N33	W69	3B
59	30/06/91	0	14	2	Max	1539E	1549D	1543	N28	E90	SF
60	07/07/91	5	7	1	Max	0120E	206	131	N26	E03	2B
61	25/08/91	15	6	2	Max	0031E	0503D	49	N24	E77	2B
62	30/09/91	6	3	3	Max	1513E	1739D	1531	S21	E32	4B
63	28/10/91	6	2	1	Max	0538E	0658D	549	S13	E15	3B
64	30/10/91	7	2	1	Max	0611E	0903D	621	S08	W25	3B
65	06/02/92	22	3	1	Max	928	1432	1028	S13	W10	2B
66	27/02/92	12	2	3	Max	922	1303	1256	N06	W02	3B
67	07/03/92	14	5	3	Max	NO FLARE ASSOCIATION					
68	15/03/92	9	3	3	Max	121	350	146	S14	E29	3B
69	09/05/92	0	4	2	Max	1537E	1856D	1540U	S26	E08	4B
70	25/06/92	21	7	1	Max	1947	2229	2011U	N09	W67	2B
71	04/08/92	15	4	1	Max	629	825	702	S09	E68	1N
72	30/10/92	19	10	1	Max	1702	2203	1730	S22	W61	2B
73	23/11/92	19	2	3	Max	2018	2101D	2029U	S08	W89	SF
74	29/11/92	11	2	3	Max	1230	1250	1233	S16	W01	1N
75	04/03/93	13	3	3	Max	1016	1106	1017	N10	W23	1N
76	06/03/93	22	4	3	Max	2014	2218	2033	S04	E29	3B
77	12/03/93	19	3	1	Max	1648	2042	1820	S00	W51	3B
78	15/05/93	0	2	3	Max	2154	2458D	2253	N19	W48	2B
79	07/06/93	16	2	1	Max	1354	1712	1422	S10	W30	2B
80	24/06/93	20	5	3+2	Max	719	0946D	0741U	S11	E64	2B
81	20/02/94	2	3	1	Max	138	138	308	N09	W02	3B
82	20/04/98	12	7	1	Max	938	1118	1021			
83	30/04/98	2	2	3	Max	1606	1659	1637	S18	E20	
84	02/05/98	14	3	1	Max	1331	1351	1342	S15	W15	3B
85	06/05/98	8	2	1	Max	758	820	809	S11	W65	1N
86	09/05/98	6	3	1	Max	304	355	340			
87	16/06/98	4	3	3	Max	1803	1928	1842			
88	22/08/98	10	3	3	Max	2357	16	9	N42	E51	2B
89	24/08/98	23	8	1	Max	2150	2235	2212	N35	E09	3B
90	24/09/98	19	2	1	Max	640	731	713	N18	E09	3B
91	30/09/98	15	5	1	Max	1300	1345	1900			
92	19/10/98	0	2	3	Max	NO FLARE ASSOCIATION					
93	06/11/98	6	3	1	Max	1900	2012	1955	N22	W18	
94	14/11/98	7	4	1	Max	2055	2102	2059	N19	W59	SF
95	20/01/99	22	6	1	Max	1906	2100	2004			
96	24/04/99	16	3	1	Max	NO FLARE ASSOCIATION					
97	04/05/99	0	4	2	Max	536	632	602	N15	E32	2N
98	01/06/99	22	7	1	Max	NO FLARE ASSOCIATION					

Appendix C

Impulsive and Gradual X-Ray Flare Associations for SPEs

GOES Date	Class of X-ray Flare
32141.00	LDE
02/01/88	LDE
25/03/88	Impulsive
30/06/88	Impulsive
24/08/88	LDE
04/10/88	Impulsive
12/10/88	Impulsive
08/11/88	LDE
14/11/88	Impulsive
14/12/88	Impulsive
27/12/88	Impulsive
04/01/89	LDE
18/01/89	LDE
07/03/89	LDE
17/03/89	Impulsive
23/03/89	LDE
10/04/89	LDE
22/04/89	Impulsive
01/05/89	Impulsive
20/05/89	LDE
18/06/89	Impulsive
20/06/89	Impulsive
29/06/89	Impulsive
29/06/89	LDE
25/07/89	impulsive
12/08/89	LDE
12/09/89	LDE
16/09/89	LDE
29/09/89	LDE
19/10/89	LDE
15/11/89	impulsive
18/11/89	LDE
26/11/89	LDE
30/11/89	LDE
03/02/90	impulsive
19/03/90	LDE
28/03/90	LDE
07/04/90	no flare
16/04/90	LDE
28/04/90	no flare
07/05/90	impulsive
15/05/90	LDE
21/05/90	impulsive
24/05/90	impulsive
12/06/90	LDE
13/07/90	LDE
26/07/90	LDE
31/07/90	LDE
13/08/90	impulsive
25/08/90	no flare
23/10/90	no flare

GOES Date	Class of X-ray Flare
11/11/90	LDE
23/12/90	impulsive
25/01/91	LDE
08/02/91	LDE
25/02/91	LDE
23/03/91	LDE
31/03/91	LDE
03/04/91	impulsive
10/05/91	LDE
13/05/91	LDE
18/05/91	LDE
31/05/91	impulsive
15/06/91	LDE
30/06/91	LDE
07/07/91	LDE
25/08/91	LDE
30/09/91	LDE
28/10/91	impulsive
30/10/91	LDE
06/02/92	LDE
27/02/92	LDE
07/03/92	no flare
15/03/92	LDE
09/05/92	LDE
25/06/92	LDE
04/08/92	impulsive
30/10/92	LDE
23/11/92	LDE
29/11/92	LDE
04/03/93	impulsive
06/03/93	LDE
12/03/93	LDE
20/04/98	LDE
02/05/98	LDE
06/05/98	LDE
24/08/98	LDE
24/09/98	no flare
30/09/98	LDE
06/11/98	LDE
14/11/98	impulsive
20/01/99	LDE
24/04/99	no flare
04/05/99	LDE
27/05/99	impulsive
01/06/99	impulsive
11/06/99	impulsive

Appendix D
Listing of Quiet Periods

ID	Date	Time
1	05/12/77	0
2	06/12/77	0
3	10/12/77	0
4	14/12/77	0
5	22/12/77	0
6	30/01/78	0
7	01/03/78	0
8	04/03/78	0
9	07/03/78	0
10	08/03/78	0
11	10/03/78	0
12	15/03/78	0
13	25/03/78	0
14	01/04/78	0
15	27/05/78	0
16	28/05/78	0
17	12/06/78	0
18	15/06/78	0
19	20/06/78	0
20	01/08/78	0
21	30/09/78	0
22	31/10/78	0
23	04/11/78	0
24	03/12/78	0
25	28/12/78	0
26	30/12/78	0
27	01/01/79	0
28	02/01/79	0
29	12/01/79	0
30	14/01/79	0
31	18/01/79	0
32	30/01/79	0
33	28/03/79	0
34	18/04/79	0
35	25/04/79	0
36	13/05/79	0
37	21/05/79	0
38	29/05/79	0
39	02/06/79	0
40	24/06/79	0
41	25/06/79	0
42	20/07/79	0
43	25/07/79	0
44	14/08/79	0
45	16/10/79	0
46	21/10/79	0
47	24/10/79	0
48	09/11/79	0
49	11/11/79	0
50	14/11/79	0
51	04/12/79	0
52	21/12/79	0
53	21/01/80	0
54	02/02/80	0
55	26/02/80	0
56	28/02/80	0
57	05/03/80	0
58	19/03/80	0
59	21/03/80	0
60	22/03/80	0
61	31/03/80	0
62	15/04/80	0
63	18/04/80	0
64	30/04/80	0
65	06/05/80	0
66	29/05/80	0
67	30/05/80	0
68	02/06/80	0
69	23/06/80	0
70	29/08/80	0
71	01/09/80	0
72	20/09/80	0
73	23/09/80	0
74	02/10/80	0
75	04/10/80	0
76	06/10/80	0
77	11/10/80	0
78	12/10/80	0
79	05/11/80	0
80	11/11/80	0
81	02/01/81	0
82	07/02/81	0
83	14/06/81	0
84	21/06/81	0
85	23/06/81	0
86	26/06/81	0
87	27/06/81	0
88	03/07/81	0
89	06/07/81	0
90	07/07/81	0
91	09/07/81	0
92	15/07/81	0
93	26/08/81	0
94	01/09/81	0
95	17/09/81	0
96	19/02/82	0
97	23/03/82	0
98	15/04/82	0
99	16/04/82	0
100	21/04/82	0
101	27/04/82	0
102	28/04/82	0
103	02/06/82	0
104	05/06/82	0
105	15/08/82	0
106	19/08/82	0
107	29/08/82	0
108	12/10/82	0
109	15/10/82	0
110	16/10/82	0
111	11/11/82	0
112	12/11/82	0
113	11/01/83	0
114	17/01/83	0
115	19/02/83	0
116	25/02/83	0
117	03/03/83	0
118	15/03/83	0
119	22/03/83	0
120	23/03/83	0
121	26/03/83	0
122	27/03/83	0
123	29/03/83	0
124	30/03/83	0
125	09/04/83	0
126	11/04/83	0
127	03/05/83	0
128	19/05/83	0
129	08/06/83	0
130	30/06/83	0
131	04/07/83	0
132	10/07/83	0
133	12/07/83	0
134	16/07/83	0
135	21/07/83	0
136	28/07/83	0
137	11/08/83	0
138	17/08/83	0
139	18/08/83	0
140	21/08/83	0
141	25/10/83	0
142	30/10/83	0
143	05/11/83	0
144	14/11/83	0
145	18/11/83	0
146	26/11/87	0
147	01/12/87	0
148	09/12/87	0
149	28/12/87	0
150	20/01/88	0
151	08/02/88	0
152	13/02/88	0
153	19/02/88	0
154	27/02/88	0
155	07/03/88	0
156	09/03/88	0
157	12/03/88	0
158	18/03/88	0
159	21/03/88	0

ID	Date	Time
160	10/04/88	0
161	13/04/88	0
162	22/04/88	0
163	10/05/88	0
164	11/05/88	0
165	16/05/88	0
166	26/05/88	0
167	01/06/88	0
168	12/06/88	0
169	16/06/88	0
170	24/06/88	0
171	18/07/88	0
172	25/07/88	0
173	30/07/88	0
174	19/09/88	0
175	22/09/88	0
176	29/09/88	0
177	02/10/88	0
178	07/11/88	0
179	27/02/89	0
180	01/03/89	0
181	09/04/89	0
182	14/06/89	0
183	24/07/89	0
184	08/08/89	0
185	18/12/89	0
186	08/01/90	0
187	16/01/90	0
188	22/01/90	0
189	24/01/90	0
190	07/03/90	0
191	08/03/90	0
192	26/06/90	0
193	29/06/90	0
194	25/09/90	0
195	28/09/90	0
196	04/10/90	0
197	19/10/90	0
198	23/11/90	0
199	24/11/90	0
200	01/12/90	0
201	11/12/90	0
202	09/01/91	0
203	13/03/91	0
204	16/03/91	0
205	20/03/91	0
206	25/04/91	0
207	12/05/91	0
208	28/05/91	0
209	30/05/91	0
210	26/07/91	0
211	28/07/91	0
212	31/07/91	0
213	04/08/91	0
214	05/08/91	0
215	10/08/91	0
216	19/11/91	0
217	27/11/91	0
218	06/12/91	0
219	17/12/91	0
220	20/12/91	0
221	22/12/91	0
222	29/12/91	0
223	01/01/92	0
224	15/01/92	0
225	20/01/92	0
226	23/01/92	0
227	24/02/92	0
228	14/03/92	0
229	21/04/92	0
230	26/04/92	0
231	27/04/92	0
232	03/05/92	0
233	06/05/92	0
234	07/05/92	0
235	06/06/92	0
236	09/06/92	0
237	18/06/92	0
238	20/07/92	0
239	21/07/92	0
240	23/07/92	0
241	25/07/92	0
242	29/07/92	0
243	31/07/92	0
244	02/08/92	0
245	20/08/92	0
246	21/08/92	0
247	15/09/92	0
248	17/09/92	0
249	18/09/92	0
250	22/09/92	0
251	24/09/92	0
252	02/10/92	0
253	10/10/92	0
254	29/10/92	0
255	14/12/92	0
256	17/12/92	0
257	18/12/92	0
258	26/12/92	0
259	03/01/93	0
260	07/01/93	0
261	27/01/93	0
262	31/01/93	0
263	05/02/93	0
264	21/02/93	0
265	01/03/93	0
266	29/03/93	0
267	31/03/93	0
268	04/04/93	0
269	16/04/93	0
270	21/04/93	0
271	06/05/93	0
272	09/05/93	0
273	11/05/93	0
274	12/05/93	0
275	14/05/93	0
276	30/05/93	0
277	02/06/93	0
278	03/06/93	0
279	06/06/93	0
280	17/06/93	0
281	17/07/93	0
282	18/07/93	0
283	07/08/93	0
284	09/08/93	0
285	13/08/93	0
286	16/08/93	0
287	21/08/93	0
288	09/09/93	0
289	22/09/93	0
290	26/09/93	0
291	05/10/93	0
292	15/11/93	0
293	18/11/93	0
294	19/11/93	0
295	23/11/93	0
296	01/04/98	0
297	13/06/98	0
298	30/06/98	0
299	02/07/98	0
300	03/07/98	0
301	05/07/98	0
302	13/07/98	0
303	16/07/98	0
304	24/07/98	0
305	25/07/98	0
306	20/08/98	0
307	12/09/98	0
308	28/09/98	0
309	31/10/98	0
310	01/11/98	0
311	03/11/98	0
312	28/11/98	0
313	05/12/98	0
314	11/12/98	0
315	14/12/98	0

ID	Date	Time
316	23/12/98	0
317	26/12/98	0
318	18/02/99	0
319	28/02/99	0
320	17/03/99	0
321	21/03/99	0
322	04/04/99	0
323	14/04/99	0
324	17/04/99	0
325	19/04/99	0
326	19/05/99	0
327	25/06/99	0
328	27/06/99	0
329	15/07/99	0
330	28/07/99	0
331	04/08/99	0
332	18/08/99	0
333	24/08/99	0
334	27/08/99	0
335	29/08/99	0
336	30/08/99	0
337	06/09/99	0
338	09/09/99	0
339	14/09/99	0
340	23/09/99	0

Appendix E

Neuframe Learning Algorithm Parameters

The figure below shows a screen print from the Neuframe software giving the learning and momentum rates used in all MLP networks.

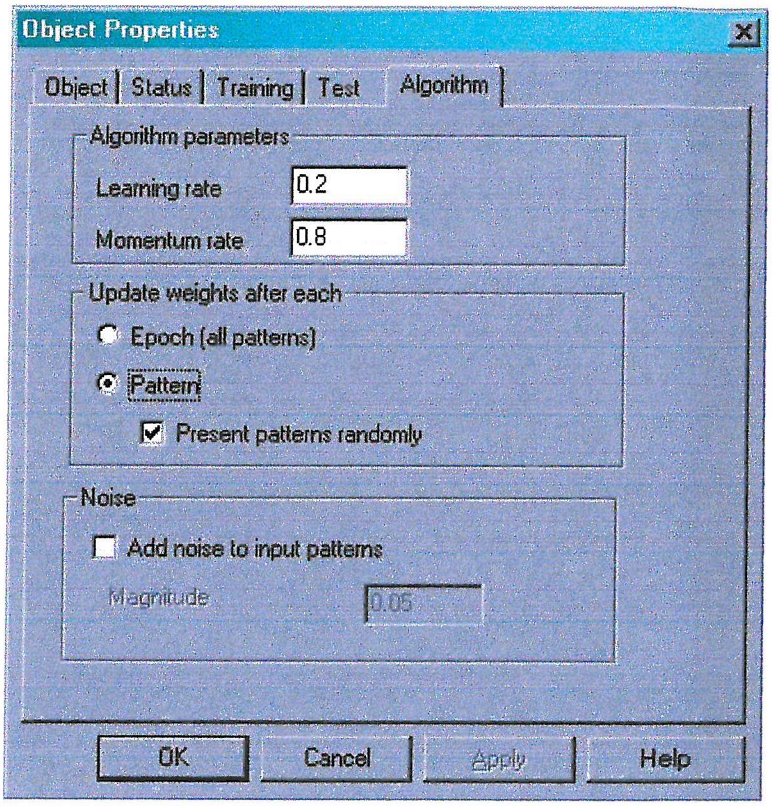

Algorithm parameters

Learning rate 0.2

Momentum rate 0.8

Update weights after each

○ Epoch (all patterns)

● Pattern

☑ Present patterns randomly

Noise

☐ Add noise to input patterns

Magnitude 0.05

During the training process the training error was monitored as a function of training epochs. Typically, the training error for any given network configuration fell rapidly to begin with and then stabilised as the number of training epochs increased. Training was defined as being complete when the training error remained approximately constant, which typically occurred between 25000 and 75000 epochs depending on the complexity of the network. An example of the training error evolution is given in the figure below and pertains to network 6 of the POPE model.

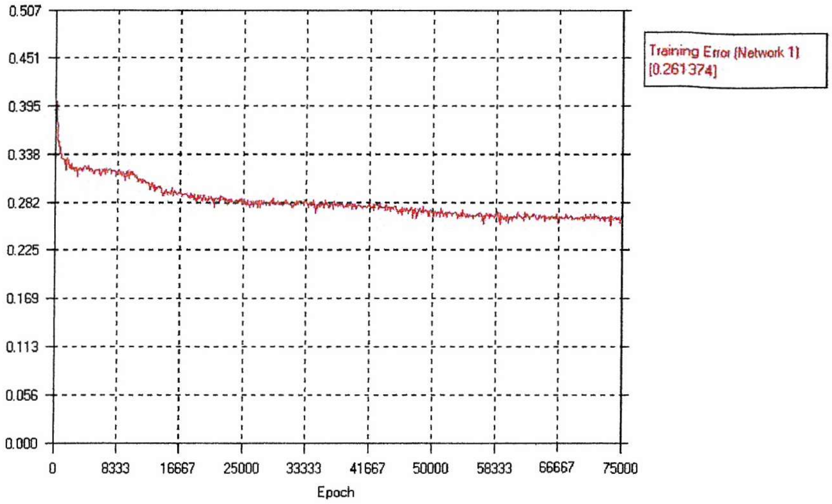

Attempts to train networks for a very large number of epochs (>75000) often resulted in the training error becoming unstable and the learning process being halted. This is thought to be due to the learning algorithm reaching the minimum point on the error surface, and then iteratively stepping around it and in so doing climbing away from the error minimum. For some networks instability was reached at <50,000 training epochs. In these cases training was reset and performed for a number of epochs just below that which triggered instability.

Appendix F
Chi-Squared Analysis Method

A Chi-Squared analysis was performed to determine the probability of a given network success rate being produced through chance alone; i.e. is a random output from the model likely to produce any of the observed success rates.

The Chi Squared test statistic χ_c is a function of the observed and expected frequencies of a distribution as shown below,

$$\chi_c^2 = \sum_{i=1}^{k} \frac{(O_i - E_i)^2}{E_i}$$

where χ_c is the computed value of a random variable having a χ^2 distribution with v degrees of freedom; O_i and E_i are the observed and expected frequencies within the ith category and k is the total number of discrete categories. Via a look-up table the computed value of χ_c can be used to find the probability that the observed distribution is the same as the expected distribution.

In this specific SPE classification problem there are 4 discrete categories, determined by the possible combinations of forecasts and observations:

1. SPE Forecast, SPE Observed.
2. SPE Forecast, Quiet Period Observed.
3. Quiet Period Forecast, Quiet Period Observed.
4. Quiet Period Forecast, SPE Observed.

Example

A chi-squared analysis was performed for a multiple model network configuration which was found to have an overall classification success rate of 61%. It was desired to know the probability of this success rate being produced through chance alone. A random prediction by the model would be expected to identify half of the SPEs and half of the Quiet Periods within the query set. The model in question was queried with a test set containing 137 SPEs and 288 Quiet Periods.

For a random output the expected distribution is:

Expected distribution from a random output		Forecast	
		SPE	QP
Observed	SPE	68.5	68.5
	QP	144	144

For the measured classification success rate of 61% the distribution was:

Measured Distribution		Forecast	
		SPE	SPE
Observed	SPE	84	53
	QP	112	176

Computing the chi-squared statistic for the two distributions gives:

$$\begin{aligned} \chi_c &= (84-68.5)^2/68.5 + (53-68.5)^2/68.5 + (112-144)^2/144 + (176-144)^2/144 \\ &= 20.57 \end{aligned}$$

Look-up tables for a distribution with 1 degree of freedom show the χ_c value to correspond to the .001 significance level; i.e. the observed success rate has a less than 0.1% probability of occurring if the output from the model were random.

Appendix G

POPE Software Specification Document

University
of Southampton

QinetiQ

Software Specification Document

Title: Predictor Of Proton Events (POPE) Software Specification Document

DERA Contract No: CU009-0000001987

Prepared by: Gareth Patrick, *University of Southampton, Astronautics Research Group*

Date: November 2001

Abstract:

This document identifies the key software requirements for a real-time solar proton event prediction model that will be developed from neural network techniques investigated under DERA contract CU009-0000001987 by the University of Southampton Astronautics Research Group. Included in the document is a description of the process to be coded, and descriptions of the components that will constitute the software.

TABLE OF CONTENTS

1 Introduction

1.1 Contractual

This document has been constructed under Work Package 4 (WP4) of DERA Contract CU009-0000001987 (Neural Network Prediction of Solar Energetic Particle Events) by The University of Southampton Astronautics Research Group.

1.2 Purpose of document

This document describes the software requirements and architectural design of a real-time solar proton prediction model "Predictor of Proton Events" (POPE) generated from the neural networks optimised in previous work [1] of the same contract.

1.3 Definitions, acronyms and abbreviations

ESA	European Space Agency
GOES	Geosynchronous Observational Environmental Satellites
JDK	Java Development Kit
MLP	Multi Layer Perceptron
NGDC	National Geophysical Data Centre
PCA	Principal Components Analysis
POPE	Predictor Of Proton Events
URL	Universal Resource Locator
UTC	Universal Coordinated Time

1.4 References

[1] Report on Work Package 3, Comparison of Non-Linear Models for the Prediction of Solar Proton Events and GOES >2MeV Trapped Electron Flux, *Gareth Patrick, Steve Gabriel.*

[2] Java How to Program, Third Edition, *Deitel &Deitel,* published by Prentice Hall

[3] GOES x-ray data site: http://www.sec.noaa.gov/ftpmenu/lists/xray.html

1.5 Overview of the document

This document describes the key requirements of the software, a description of the process to be implemented, the top-level design of the software and a brief description of the modules that will compose the software.

2 Model description

POPE is based on previous investigations which developed MLP neural networks to predict the occurrence of solar proton events using solar x-rays from GOES satellites as inputs. Given x-ray input data networks will predict 'yes' or 'no' as to the occurrence of a proton event in 48-hours time.

POPE will access real-time GOES satellite data (available from the NGDC ftp site[3]) and process it for input to a neural network model developed in WP3. The output from the neural network will be interpreted and recorded with the potential of making it available as a space weather forecasting service.

Earlier work showed that a greater potential for success was achieved when several differently trained networks of the same type were connected in parallel and fed the same query vector as an input. This creates several outputs for the same input, and allows the number of 'yes' and 'no' predictions to be counted separately and interpreted.

Figure 1 shows a schematic diagram of the model structure. X-ray data is obtained from the NGDC and detrended for the long term solar cycle (based on a polynomial calculated from historical data). The detrended input vector is fed to 10 MLP networks. In each case the input vector is projected onto principal components, using coefficients derived from each of the 10 different training that sets were used to develop the MLP networks. The inputs are then scaled using scales derived from the training sets and fed as inputs to the MLP networks. Outputs are then de-scaled and interpreted as being either a 'yes' or a 'no'. The number of yes and no predictions are counted and recorded, as is the time of the model-run and the time for which the prediction is valid.

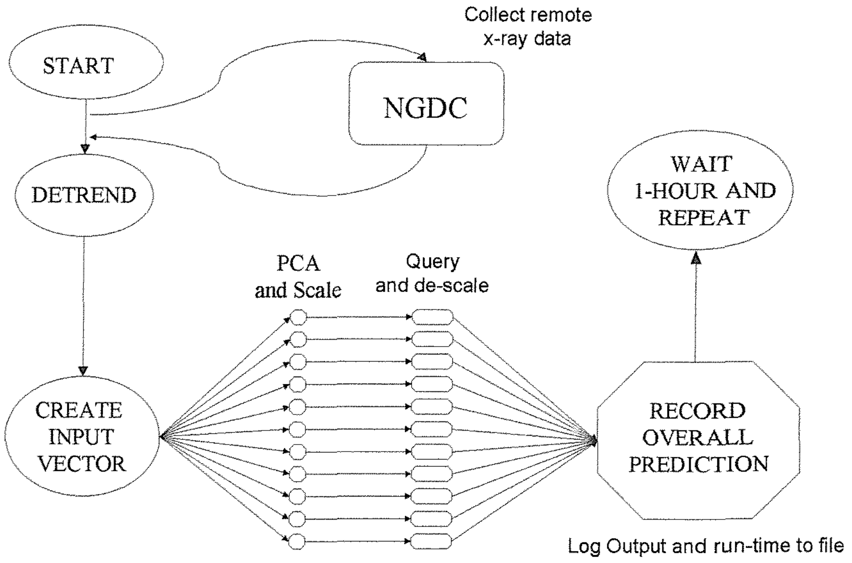

Figure 1 Schematic Diagram of POPE model

3 Specific requirements

3.1 Functional requirements

This section contains the software requirements that have been defined in order to establish the scope of the project and identify the key characteristics of the POPE software.

Requirement ID	*Description*
SR1	Model must use real-time GOES x-ray data
SR2	MLP neural network models developed in WP-3 must be integrated into the software
SR3	Model should be able to run as a stand-alone application
SR4	Model inputs and outputs should be recorded in output files to enable validation and analysis of network behaviour
SR5	Software should have the ability to run in an automated way
SR6	It should be possible to integrate other future models into the software with minimal adjustment to the code. (i.e. it should be possible to 'plug-in' new models to a real-time framework).

Table 1 Software Requirements

3.2 Performance requirements

Each prediction run should be able to execute sufficiently fast so that predictions are timely. The 48-hour prediction lead time means that practically there is no performance requirement.

3.3 Interface requirements

There is no requirement to interface with POPE once it has begun execution. Key variables concerning network configuration can be altered in a configuration text file. Other variables within the software will be set in the source code before the program is compiled and run. Unless adding other prediction models or changing the source of input data it should not be necessary to alter the source code.

3.4 Operational requirements

POPE will need to be running on a machine which has a constant internet connection.

3.5 Resource requirements

In order to develop source code for MLP networks it is necessary to extract code from the commercial neural network package "Neuframe" version 4. This is not required to run POPE, but will be used to produce source code for specific MLP networks.

3.6 Verification requirements

Code shall be verified during the construction stage by comparing source code MLP network outputs with the same networks in Neuframe.

3.7 Acceptance testing requirements

None. The software is intended as a validation project within itself.

3.8 Documentation requirements

Software User Manual and Software Specification Document.

3.9 Portability requirements

Code will be developed on a PC running windows 98, although it is desirable for POPE to be able to run under Windows and UNIX environments.

3.10 Quality Requirements

Code shall be written and documented in accordance with ESA software engineering standards for small projects (ESA PPS-05-0).

3.11 Reliability requirements

There is no significant consequence resulting from the program crashing. Reliability shall be made as high as possible within the time frame allocated.

3.12 Maintainability requirements

Code shall be written in JAVA (version 1.3.01). There will be no use of deprecated functions and detailed information will be contained in the source code relating to areas of code that have potential for modification.

4 System Design

4.1 Design Method

POPE will be coded in JAVA using an object oriented approach as described in "JAVA How to Program", Deitel and Deitel. JAVA facilitates the use of remote data and is platform independent.

4.2 Decomposition Description

Figure 2 below outlines the software structure that will be used to achieve the stated requirements. Six software classes will be developed for specific stages of the process.

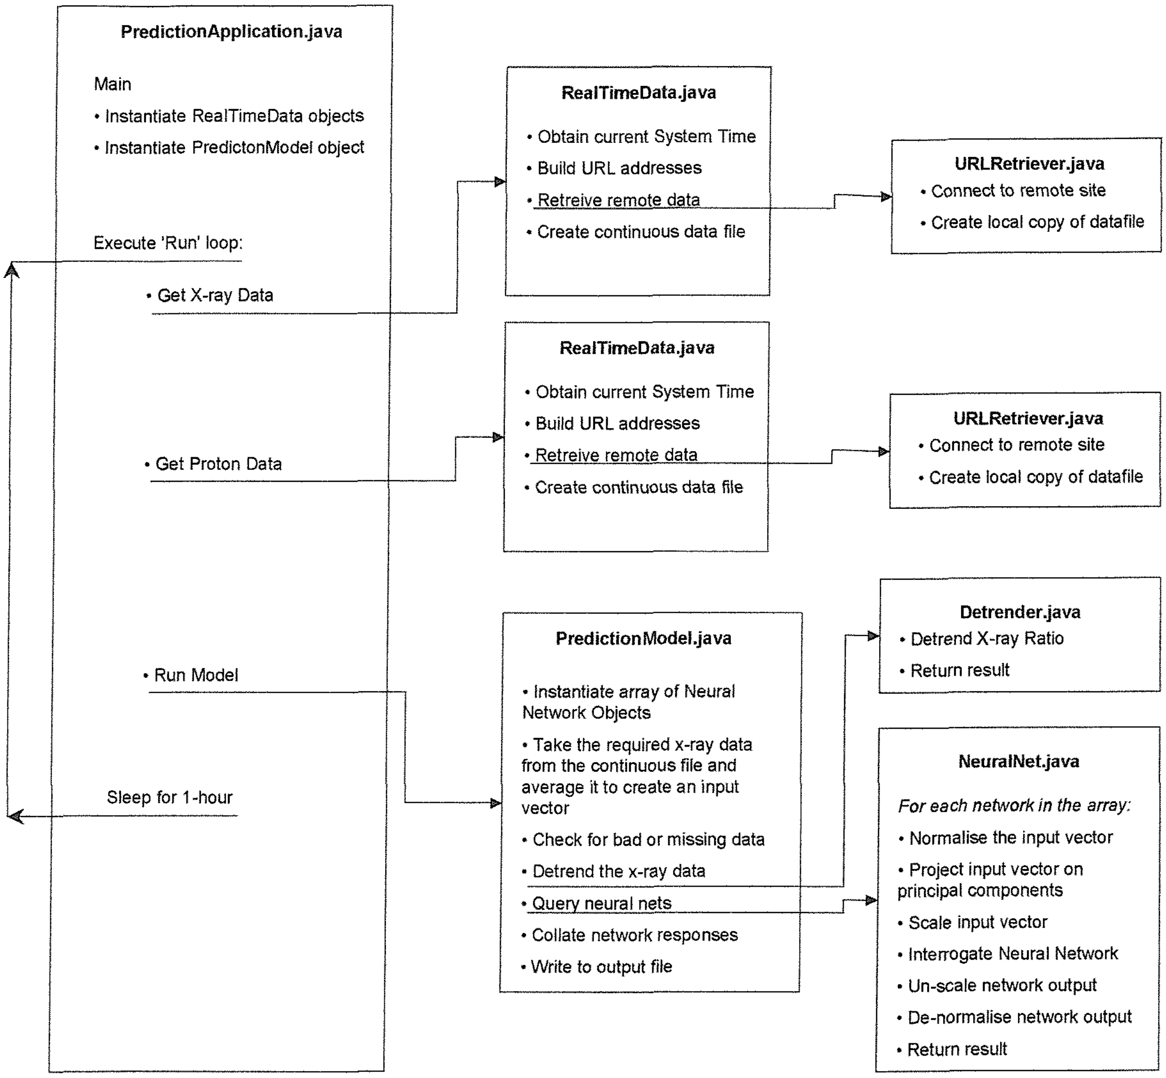

Figure 2 Software Architecture Diagram

PredictionApplication.java contains the main function which will instantiate a PredictionApplication object. The PredictionApplication object will create two instances of the RealTimeData object (one for accessing GOES x-ray data and one for GOES proton data), and one instance of a PredictionModel object.

A looping thread is started within the PredictionModel object which will perform the tasks of retrieving remote data, querying the neural models and writing output to a data file. The thread will then enter a sleep mode for 1-hour, after which it will begin the sequence again, unless halted by terminating the program.

The RealTimeData class will deal specifically with obtaining real-time GOES data. It will contain methods to calculate the current UTC time from the system time, and will generate filenames for GOES 5-min x-ray or proton data using the current UTC time. Retrieval of the remote data will be performed via a URLRetriever object. The URLRetriever class contacts a given URL address and creates a local copy of the file that is found there. GOES data files from several days will be collated into one file.

Although the models use only x-ray data as an input, real time data will be retrieved for protons also. This will be saved to an output file to allow a ready comparison between the model output and the >10MeV proton flux.

Once real-time data has been retrieved the prediction model will be run. The PredictionModel class will contain instance variables relating to the configuration of the model, such as the length of delay vector required, the time average of delay vector required and the lead-time. Based on these variables the raw delay vector is created from the continuous GOES data file.

Instances of NeuralNet objects are created within the PredictionModel object. This allows an array of different NeuralNet objects to be created for the same model characteristics. It is possible to just use one neural network, but by creating an array of 10 NeuralNet objects a multiple model configuration will be used, whereby 10 differently trained networks can be queried with the same input vector.

The NeuralNet class will contain variables and methods specific to each MLP neural network, such as weight vectors, training set characteristics and PCA coefficients. Each NeuralNet object will receive the input vector created in the PredictionModel object, then normalise, project and scale the vector before sending it as an input to the MLP network itself. The output shall be un-scaled and de-normalised before being returned to the PredictionModel object.

The PredictionModel object will produce and update an output file containing the model outputs, time, input vector and recent GOES x-ray and proton fluxes. This file can be read into other spreadsheet software for analysis.

After running the PredictionModel the thread will enter a sleep model for 1-hour before repeating the process and generating a new prediction. If desired, the source code can be altered by the user to use a sleep time different to 1-hour.

5 Component Description

5.1 PredictionApplication.java

5.1.1 Type

Class

5.1.2 Function

The PredictionApplication class will contain the 'main' method and the thread of top-level program execution calls to get remote real-time data, query the network models and create an output file. The thread will be placed in a repeating loop which sleeps for 1-hour after each execution. This class controls the top-level program sequence.

5.1.3 Interfaces

The program will be started from a command line in the operating system being used (MS-DOS in windows). A commentary will be written to screen as the statements in the run-thread relating to the main stages of the software process are executed.

The run-thread will call RealTimeData methods in order to retrieve data for the current UTC time. The run-thread will call PredictionModel methods in order to produce an input vector from the real-time data and query the MLP networks

5.1.4 Dependencies

For the run-thread to finish executing and reach the sleep mode all components of the process must have completed successfully.

5.1.5 Processing

The main method will instantiate a PredictionApplication object and invoke its run method. Real time GOES proton and x-ray data will be retrieved via RealTimeData objects inside the run method, and a PredictionModel object will be used to query MLP networks with the real-time data. This will end the run method, and a sleep mode will ensue for 1-hour. The run-method will then repeat. This will continue indefinitely until the program is terminated by killing the process from the operating system.

5.1.6 Data

The PredictionApplication object contains two instances of a RealTimeData object and one instance of a PredictionModel object. No actual data processing is performed with PredictionApplication though - it acts only to control the processing sequence and repeating loop.

5.1.7 Resources

None

5.2 RealTimeData.java

5.2.1 Type

Class

5.2.2 Function

Will create URL addresses and retrieve several days worth of GOES x-ray or proton data relative to the current system time and collate it into a continuous file (of 5 minute averages) ordering the data so that the most recent is at the top of the file. The class will also contain any other methods relating to the compilation of averages from recent GOES data.

5.2.3 Interfaces

RealTimeData methods will be called from the run-thread in the PredictionApplication class

Methods of RealTimeData will read local copies of GOES x-ray and proton data files to collate into one file. The collated file will be saved as a local file.

URL addresses will be passed to methods of a URLRetriever object for retrieval.

5.2.4 Dependencies

Some Methods of the class will require local copies of the GOES real-time data files. These files will need to be retrieved and of the expected format.

5.2.5 Processing

RealTimeData will calculate the current UTC time from the system time.

RealTimeData will be able to create file name strings and URL addresses from the UTC time.

RealTimeData will be able to calculate recent flux averages from retrieved GOES data for protons and x-rays.

RealTimeData will be able to collate several 1-day GOES data files into one file to provide a continuous record.

5.2.6 Data

RealTimeData will require the exact file stems relating to the GOES data files (e.g. G8_xr5m.txt) and the number of days for which data should be retrieved for, relative to the current UTC time. The exact file format of GOES data will also be required in order to read data correctly.

5.2.7 Resources

Internet Connection for remote file retrieval

5.3 URLRetriever.java

5.3.1 Type

class

5.3.2 Function

Create a local copy of a URL address.

5.3.3 Interfaces

A URL address will be passed to a URLRetriever method. The file at a URL address will be saved as a local copy.

A configuration file must be read containing the correct proxies for use of a remote connection.

5.3.4 Dependencies

Internet connection proxies must be set correctly.

5.3.5 Processing

The system properties will need to be modified to use proxy internet settings (host and port).

5.3.6 Data

URLRetriever will require proxy internet settings to modify system properties.

URLRetriever will require the local directory and filename to save a URL to.

5.3.7 Resources

Internet connection.

5.4 PredictionModel.java

5.4.1 Type

Class

5.4.2 Function

This class will contain all variables and methods for a specific prediction model configuration, such as size and averaging period of input vector and lead-time. This class will contain methods to create an input vector from a file of continuous GOES data and will send this input vector as a query to a neural network.

5.4.3 Interfaces

Methods from the class will read a continuous file of GOES x-ray data in order to create the desired input vector.

The input vector will be passed to a Detrender method and a detrended input vector received.

A detrended input vector will be passed to NeuralNet objects as a query.

PredictionModel will receive the output from NeuralNet objects.

A PredictionModel method will write data to an output file.

5.4.4 Dependencies

A continuous file of GOES 5-minute x-ray data must have been created by a RealTimeData object.

A detrender object must be instantiated with correct detrending coefficients.

5.4.5 Processing

When creating the input delay vector 5-minute averages will be converted into 3-hour averages. Checks will be made for bad or missing data and the number of non-valid data points recorded.

Within the PredictionModel object an array of NeuralNet objects will be instantiated. For each NeuralNet object weights, training set statistics and scaling coefficients will need to be initialised.

A header shall be written to the output file, which will contain the UTC date, time and Julian Day at which the model ran, the UTC date, time and Julian Day of the model prediction, a copy of the raw x-ray input vector, the number of missing data values in the 5-minute x-ray data, the numerical output of each MLP network, the number of networks predicting an event in 48-hours time and the previous 1-hours averages of GOES x-ray and >10MeV proton fluxes.

The results file will be appended each time the run-thread completes a cycle.

5.4.6 Data

The PredictionModel object will require the lead-time, delay vector length and input average for the desired configuration. (In this case, these are 48 hours, 72 hours and 3 hours respectively).

The weights matrices, training set statistics (average and standard deviation of each input dimension) and the scaling coefficients will be required for each NeuralNet object instantiated within PredictionModel.

5.4.7 Resources

None

5.5 Detrender.java

5.5.1 Type

Class

5.5.2 Function

To detrend the GOES x-ray ratio data by removing the long term solar cycle variation.

5.5.3 Interfaces

Detrender.java will receive 5-minute x-ray data and will return the detrended x-ray time series.

5.5.4 Dependencies

Coefficients defining the detrending function (a second order polynomial) will need to be set.

5.5.5 Processing

Detrender will receive a raw x-ray ratio data as a time series, and calculate the trend function at each point along the time series. The trend will then be subtracted from the raw values to generate a detrended time series.

5.5.6 Data

Detrender will require the coefficients of the 2nd order polynomial function.

It is envisioned that the detrending function for the solar cycle will be calculated (externally) by plotting the x-ray ratio as a function of Julian Day and then fitting a 2nd order polynomial using least squares. To ensure accuracy of the coefficients the Julian day will need to be standardised by subtracting the average and dividing by the standard deviation. These values of the average and standard deviation will be needed by the class so that the current Julian Day can be standardised in the same way when calculating the detrending function.

5.5.7 Resources

None

5.6 NeuralNet.java

5.6.1 Type

Class

5.6.2 Function

To feed an input vector to a MLP neural network and obtain an output. This class will contain all the variables and methods that define a single MLP network, such as weights matrices, scaling coefficients and number of layers and nodes.

5.6.3 Interfaces

Each NeuralNet object will receive the detrended input vector created by a PredictionModel object. An output will be returned to the PredictionModel object.

A file will be read by a NeuralNet object containing the Principal Component Analysis coefficients.

5.6.4 Dependencies

Before a NeuralNet object can be queried values for weights and scaling coefficients will need to be set.

If a principal components analysis is to be performed then a file containing its coefficients in the correct format needs to be present.

5.6.5 Processing

The NeuralNet object will take the input vector created by a PredictionModel object and perform a normalisation and a PCA projection. The correct number of principal components will be selected to form the final input vector, and this will be scaled using coefficients derived from the original training set. The scaled input vector will be passed through a MLP network structure with weights set and the output un-scaled and de-normalised.

5.6.6 Data

Each NeuralNet object will require a weights matrix, the number of layers and number of nodes per layer in the MLP network, the average and standard deviation of each dimension in the original training set (to allow normalisation), the PCA projection coefficients and the coefficients of the input scale that map inputs to values of between 0 and 1.

5.6.7 Resources

None

6 Feasibility and Resource Estimates

Software will be built using the Java Development Kit version 1.3.1_01 available free from Sun Microsystems. Development will be carried out on a PC running Windows 98 and Borland Jbuilder (University Edition). Neural networks will be implemented by extracting code from Neuframe version 4.

The minimal requirement to maintain the code will be an installation of JDK 1.3.1_01 and the ability to edit text files, although a java development package will facilitate code editing.

In order to run, the software will require an MS-DOS prompt via a Windows operating system and a constant internet connection. A server type machine would provide a suitable platform. Although developed in Windows it is expected that POPE will be able to run on a UNIX/Linux platform as well.

7 Software Requirements vs Components Traceability Matrix

	SR1	SR2	SR3	SR4	SR5	SR6
PredictionApplication.java			✓		✓	
RealTimeData.java	✓					
URLRetriever.java	✓					
PredictionModel.java				✓		✓
Detrender.java						✓
NeuralNet.java		✓				✓

Table 2 Software Requirements vs. Components Traceability Matrix

Appendix H
POPE Software User Manual

POPE v1.0
Software User Manual

Predictor Of Proton Events (POPE) Software User Manual

February 2002

DERA contract CU009-0000001987

G. Patrick
University of Southampton
Astronautics Research Group

TABLE OF CONTENTS

1 Introduction

1.1 Intended Readership

The Software User Manual should be read by anyone intending to install, use or modify POPE, or for anyone wishing to understand the process supported by POPE.

1.2 Applicability Statement

This SUM applies to version one of POPE, completed February 5th 2002.

1.3 Acronyms

SUM	Software User Manual
POPE	Predictor of Proton Events
NGDC	National Geophysical Data Centre
MLP	Multi Layer Perceptron
SPE	Solar Proton Event
GOES	Geosynchronous Observational Environment Satellite
UTC	Universally Coordinated Time

1.4 References

[1] Work Package 3 Report "Comparison of Non-Linear Models for the Prediction of Solar Proton Events and GOES >2MeV Trapped Electron Flux ", Gareth Patrick, July 2001.

1.5 Purpose

This document gives an overview of the process that is supported by POPE and describes in detail how the code operates. It also contains instructions on how to install and run POPE and gives explanations of the output files that are produced by POPE. In addition, reference information is given concerning the source code and how to implement further prediction models within the existing software.

POPE was written in order to validate prediction models created under previous work packages of the contract. The software implements MLP neural networks into a real-time framework in order to create a real-time SPE prediction model. POPE obtains GOES satellite x-ray data from a remote ftp site and processes it for input to a prediction model. Neural networks are then queried and predict whether or not a SPE will occur in 48-hours time. Results are written to an output file.

POPE has been designed so that additional models can be coded and plugged into the real-time framework with minimum effort, and can serve as a validation platform for any type of model using GOES satellite data as inputs.

1.6 How to use this document

Section 2 provides a description of the software architecture and process, and should be read by a user wishing to understand how POPE operates and the individual tasks it actually performs.

Specific installation and running instructions are given in sections 3 and 4 and should be read prior to installation. Section 4 also includes problems which may be encountered when attempting to run the software.

Section 5 is a reference section and is intended for developers wishing to modify the existing code. Instructions are given on how to create and implement neural networks from Neuframe and how the software has been designed so as to allow different models to be plugged into the real-time interface.

A detailed description of each class and method is not given in this document, but is available in the form of HTML documentation files created from the source files using the *javadoc* facility. These files are contained within the 'documentation' folder. Extensive comments are given in the source code and the javadoc files, and only an overview of the classes will be given in this document.

1.7 Related Documents

Software Specification Manual.

2 POPE Overview

2.1 Process Overview

POPE has been written to implement in real-time MLP neural network models that were developed in Work Package 3 of the contract [1]. These models use the ratio of the two GOES x-ray channels as an input over a 72-hour period, and require access to several days worth of real-time GOES x-ray data. Figure 1 below shows a schematic diagram of the basic process to be supported.

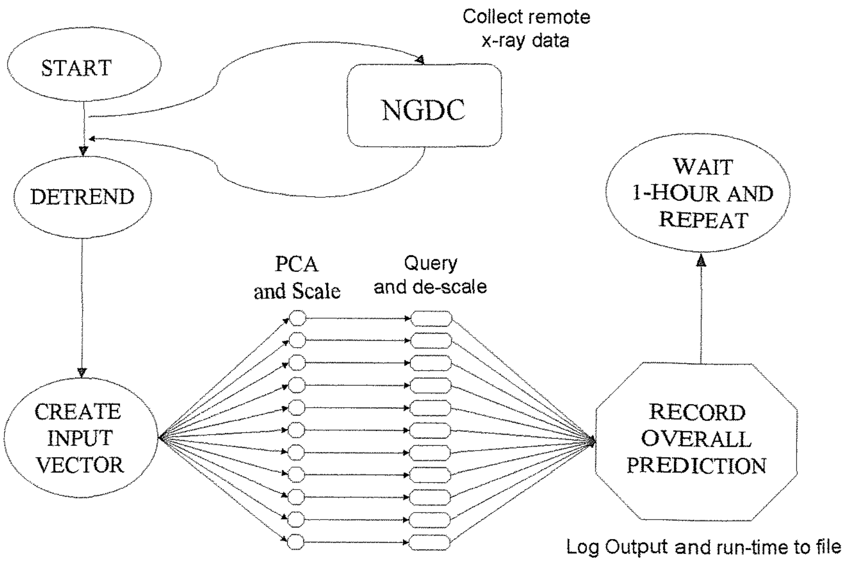

Figure 1 Schematic diagram of the basic process to be supported.

Once data has been obtained it must be pre-processed in a number of stages and then fed to the neural networks. Each neural network will generate an output and predict whether or not a SPE will occur in 48-hours time. A prediction will be generated once every 60 minutes using the latest GOES satellite data. An array of 10 differently trained neural networks will be fed the same input vector to create 10 different responses to the same query. Development work showed that better performance was generally achieved when the outputs from several networks were combined in order to produce a 'majority vote' prediction.

2.2 Description of neural network prediction model

MLP networks were created under WP-3 using a training set of 60 events and 60 quiet periods. The aim was to train a network to differentiate between the case of ‘event’ and ‘quiet period’ by using x-ray data from well before the event itself. Networks were assessed using unseen data and were found to predict the outcome correctly in 65% of cases. Individual network performance measurements are given in Appendix A.

The neural network prediction models that are implemented in this code take the ratio of the two GOES x-ray channels (XS/XL) as an input and predict ‘yes’ or ‘no’ as to whether or not an event will occur in 48-hours time.

The x-ray ratio is detrended to remove the long term solar cycle trend. This trend function is calculated by fitting a 2nd order polynomial to the x-ray ratio over the active period of the current solar cycle. The detrended x-ray ratio is calculated by subtracting the trend from the actual value of the ratio at each point in time.

The input vector to the neural network consists of a delay vector of detrended XS/XL spanning from –120 to –48 hours relative to current time, broken into 24 3-hour intervals. This is shown schematically in Figure 2.

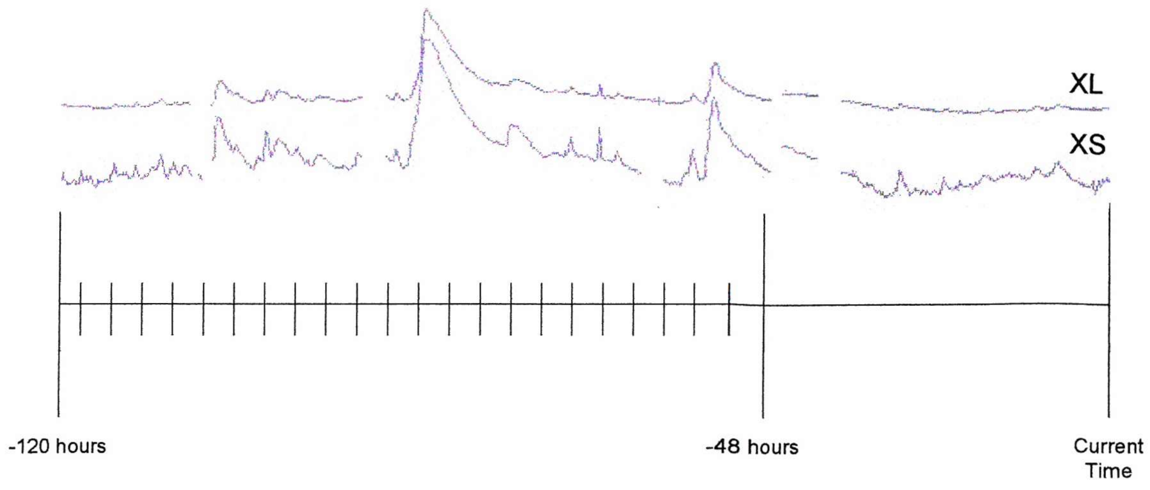

Figure 2 Input vector to the neural model

The input vector must be normalised and projected onto principal axes via a PCA technique. The statistics used for normalising the input vector and the coefficients used to perform the projection are derived from the training set that was used to build the neural model. The first six principal components are taken as the input vector to the neural network.

Each element of the input vector is scaled to a new value of between 0 and 1 using the same scale that was applied to the training set. The network is then queried and produces a numerical output which is un-scaled and de-normalised prior to interpretation. Networks have been trained using a value of 0 to represent a quiet period and a value of 100 to represent an event, thus a decision threshold of 50 is used to interpret the network output. A value of ≥50 indicates that an event will occur in 48-hours time, a value of <50 indicates that no event will occur in 48-hours time.

By feeding the same delay vector to several differently trained networks several different predictions can be generated for 1 case. Previous experience has shown that slightly higher success is achieved if the outputs from several networks are combined to form a 'majority vote' as to whether or not the outcome is an 'event' or not.

2.3 Software Overview

POPE has been written using six JAVA classes which each perform specific stages of the process. The distinct tasks carried out by POPE can be broken down as follows:

- Retrieval of real-time GOES satellite data from a remote ftp site. (Performed by URLRetriever.java)
- Creation of a data file made by stitching together several GOES data files. (Performed by RealTimeData.java)
- Detrending of x-ray data to remove the long term solar cycle trend. (Performed by Detrender.java)

- Creation of a delay vector from continuous GOES x-ray data.
  (Performed by PredictionModel.java)
- Pre-processing of the delay vector (normalising, PCA and scaling), and running a neural network model with the processed delay vector.
  (Performed by NeuralNet.java)
- Interpreting and recording the network outputs and other information relevant to each prediction.
  (Performed by PredictionModel.java)
- A control loop which repeats the prediction process every 60 minutes.
  (Performed by PredictionApplication.java)

A more detailed description of the software execution stages are given in the following section

2.4 Software Architecture

Figure 3 is a software flow-diagram displaying the interaction between classes and the specific tasks that are carried out each time a prediction process is run. Grey boxes indicate input and/or output files that are accessed or created during each run cycle.

Control of the run process is governed by PredictionApplication.java. This class contains the main function which runs a prediction once every 60 minutes via a repeating loop. Each prediction run consists of three main tasks; retrieving x-ray data, retrieving proton data and running the model.

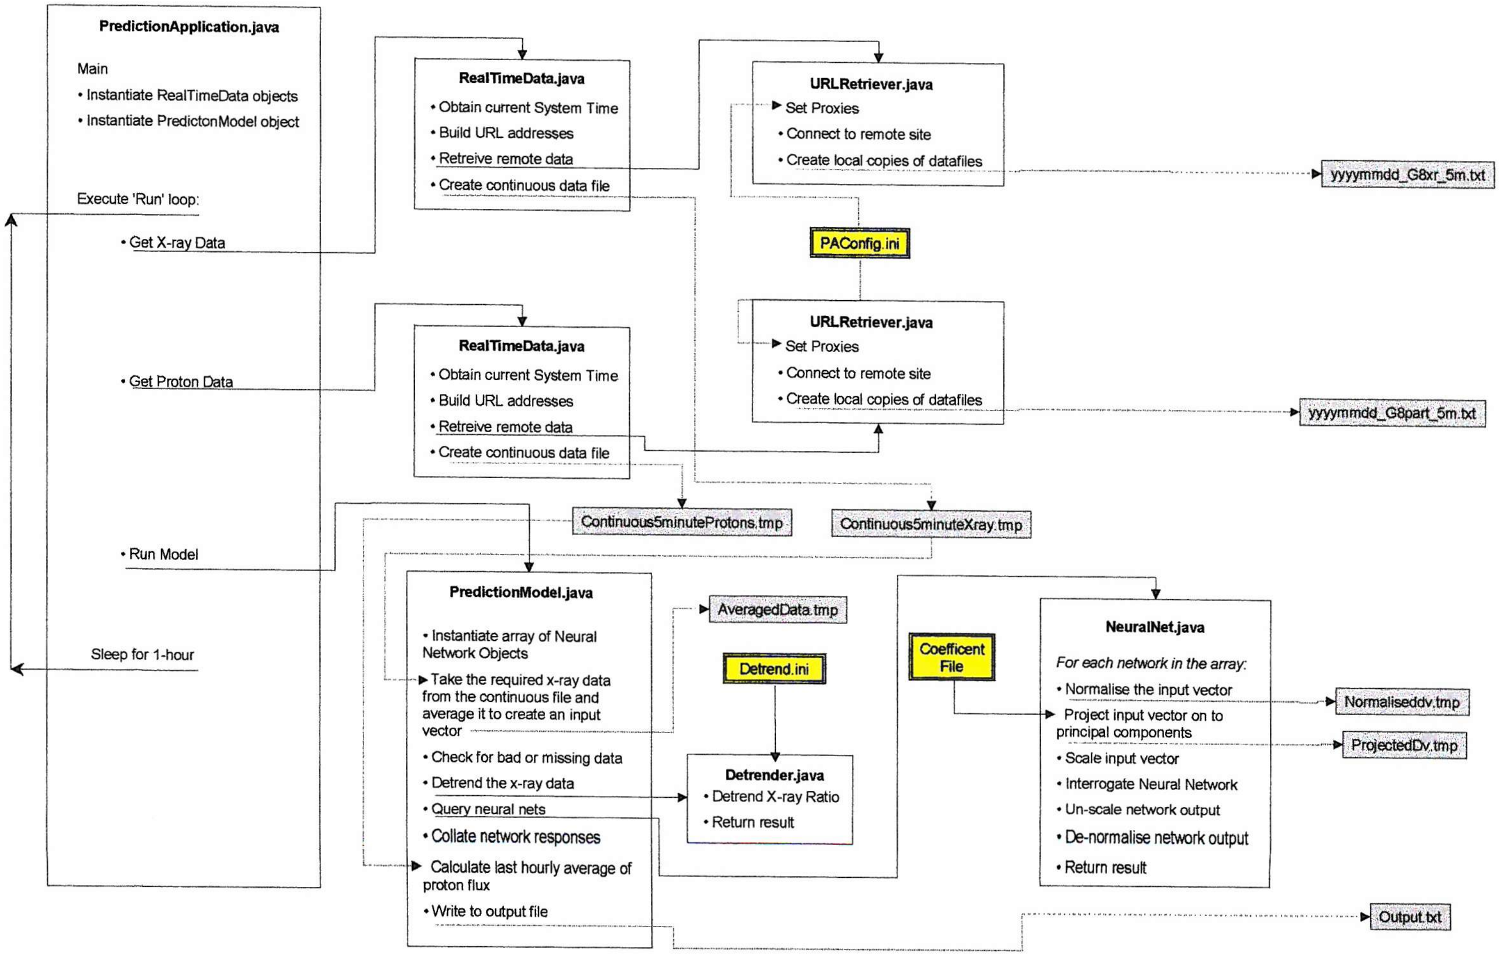

Figure 3 Schematic Diagram showing software architecture, interfacing and file outputs.

The retrieval of GOES data is performed by a RealTimeData object. This obtains the current UTC time from the system clock then creates filenames for GOES x-ray or proton data for the current day and the preceding 5 days. These URL addresses are then passed to a URLRetriever object which accesses the URL and saves the GOES files as local copies. In order to access remote data from behind a firewall it may be necessary to use proxy settings. This information is provided by the user in the PAConfig.ini file which is accessed by the URLRetriever object at run-time. Once GOES data files for the current day and previous 5-days have been retrieved the RealTimeData object removes their headers and stitches the files together to form a continuous data file with the most recent data at the top of the file. The ultimate purpose of the RealTimeData object is to create a continuous data file which can be processed further by other objects.

After retrieving GOES x-ray and proton data the prediction model is run. (Note that the prediction models implemented here use only the ratio of the GOES x-ray channels as an input, but by also retrieving proton data the current proton flux can be recorded in an output file to allow a comparison with the model output).

The prediction model is defined by the PredictionModel object. This reads in the continuous 5-minute resolution x-ray data file that was created by the RealTimeData object and detrends the XS/XL ratio using the detrending function defined within the Detrender object. The PredictionModel object then creates a delay vector from the detrended 5-minute resolution x-ray ratio, consisting of 24 3-hour averages from the 72 hour period prior to the current UTC time. The PredictionModel object serves to create the delay vector that will be passed to a neural network. For reference the delay vector is saved in a temporary file.

Contained within the PredictionModel object is an array of 10 NeuralNet objects which define 10 different MLP networks that are all queried with the same delay vector. Upon creation each NeuralNet object is initialised with a matrix of weights, the average and standard deviation of the training set on which the model was based (to allow the delay vector to be normalised) and the coefficients of the scale that is applied to the projected delay vector to produce values of between 0 and 1.

Each NeuralNet object is passed the detrended input vector created by the PredictionModel object, and normalises it according to the training set statistics. The normalised vector is then projected using PCA coefficients which are read from a coefficient file, and which have been derived from the training set. The first six principal components are taken as inputs to the network and are scaled using the scaling coefficients derived from the training set. The network output is then un-scaled and de-normalised to produce a final output which is returned to the PredictionModel object. For reference, temporary files containing the normalised delay vector and the projected delay vector are created.

The PredictionModel object collates the outputs from the 10 NeuralNet objects and records the number of networks that are predicting an event in 48-hours time. The result is printed to screen and appended to a results file with other relevant information, including current values of x-ray and proton flux.

The run process is now complete and the software enters a sleep mode for sixty minutes, after which another data retrieval and model-run will be performed.

3 Installation Instructions

3.1 Minimum Requirements

POPE must be installed on a machine which has constant internet access.

POPE can be installed under a Windows environment (via an MS-DOS prompt) or a UNIX/Linux environment.

The platform will require an installation of Java SDK version 1.3.1 or above, available from java.sun.com.

3.2 Components

The POPE software should consist of the files shown in Table 1.

Filename	Description
PredictionApplication.java	Class source file
RealTimeData.java	Class source file
URLRetriever.java	Class source file
PredictionModel.java	Class source file
Detrender.java	Class source file
NeuralNet.java	Class source file
PAConfig.ini	Holds proxy internet settings
Detrend.ini	Holds the detrending function coefficients and normalising statistics
coeff_120fq1.txt	Holds PCA coefficients for network 1
coeff_120fq2.txt	Holds PCA coefficients for network 2
coeff_120fq3.txt	Holds PCA coefficients for network 3
coeff_120fq4.txt	Holds PCA coefficients for network 4
coeff_120fq5.txt	Holds PCA coefficients for network 5
coeff_120fq6.txt	Holds PCA coefficients for network 6
coeff_120fq7.txt	Holds PCA coefficients for network 7
coeff_120fq8.txt	Holds PCA coefficients for network 8
coeff_120fq9.txt	Holds PCA coefficients for network 9
coeff_120fq10.txt	Holds PCA coefficients for network 10

Table 1 Files needed for POPE installation

3.3 Installing POPE

Place all of the POPE files into a single directory on the hard drive. If in windows, open an MS-DOS prompt and change directory to the one that contains the POPE files. If in UNIX/Linux use a command window.

To compile the software type:

```
javac PredictionApplication.java
```

3.4 Altering the PAConfig.ini file

The PAConfig.ini file contains proxy settings that will be required to access the internet from behind a firewall. Information required is the name of the proxy host and the number of the proxy port. If unknown this information can be obtained from a web browser operating on the same machine, or the system administrator. Alter the PAConfig.ini file so as to read:

```
proxyHost = name_of_proxy_host
proxyPort = number_of_proxy_port
```

If the platform being used can access the internet directly then no proxy settings are required and the PAConfig.ini file should read:

```
proxyHost =
proxyPort =
```

(i.e. the values are set to null).

At run-time the PAconfig.ini file will be accessed and the system properties changed to reflect the proxy host and proxy port information given.

4 Using POPE

4.1 Starting POPE

Open an MS-DOS prompt and change directory to the one containing the POPE files. To run the program type:

```
Java PredictionApplication
```

A message will appear stating that the POPE run thread has begun and GOES x-ray and proton data files will be downloaded from the NGDC ftp site. Messages will appear stating which files are being downloaded. Depending on the speed of the internet connection it may take several seconds to download each file.

After all files have been downloaded a message will appear stating that the prediction model is being run with real time data. The number of networks that are predicting an event will be shown along with the current UTC time and the time for which the prediction is valid. The message 'waiting for next prediction cycle' will then appear and the program will enter a sleep mode for 60 minutes before repeating the process.

Note that POPE requires no other inputs from the user. Any modification to the program (such as sleep-time for example) must be made to the source code, and then the program recompiled using the installation instructions above.

4.2 Stopping POPE

POPE has been designed to run indefinitely and will make predictions once every sixty minutes until the program is halted. To stop POPE running the process must be killed, which can be done by pressing Ctrl-C whilst the MS-DOS window is open.

Some users may prefer to disable the repeating loop in the source code and run POPE within a task-manager program. POPE will terminate after each prediction run, meaning the program will not be running continually, and the sleep interval will be controlled by the task manager. Instructions on how to disable the loop structure in the source code are given in section 5.

4.3 Output Files

Table 2 lists the files that are produced by POPE each time a prediction-run is performed. Most files simply provide a way to check that the correct pre-processing is being applied to the delay vector and have only been produced for reference purposes during development.

Filename	Description
yyyymmdd_G8xr_5m.txt	Downloaded GOES x-ray file (contained in the 'retrieved' directory).
yyyymmdd_G8part_5m.txt	Downloads GOES proton file (contained in the 'retrieved' directory).
Continuous5minuteXray.tmp	Contains continuous GOES x-ray data from current UTC time to 5 days previous.
Continuous5minuteProtons.tmp	Contains continuous GOES proton data from current UTC time to 1 day previous.
AveragedData.tmp	Contains the input vector derived from Real-time data consisting of the detrended values of XS/XL.
Normaliseddv.tmp	Contains the input vector after it has been normalised by a NeuralNet object. Note that each network in the array overwrites the Normaliseddv.tmp file each time it runs.
ProjectedDv.tmp	Contains the input vector after it has been projected using PCA coefficients of a NeuralNet object. Note that each network in the array overwrites the ProjectedDv.tmp file each time it runs.
Output.txt	Contains the Network Outputs, the time of each run and other relevant information – see below.

Table 2 Files produced by POPE during each prediction run.

Each time a prediction run is made, information is saved to the Output.txt file. The exact information recorded is given below in Table 3.

Column Heading	Example	Description
Model run time	05/02/2002 14:25	Date and UTC time of the most recent GOES data entry*
Julian Day run time	52310.601	Julian Day of the most recent GOES data entry
Prediction time	07/02/2002 14:25	Date and UTC time for which the prediction is valid for (calculated by adding the lead time to the model run time)
Julian day prediction time	52312.601	Julian Day of the time for which the prediction is valid.
DV1 – DV24	-0.760	Contains each of the 24 elements of the detrended delay vector that is passed to each network as an input.
NET1 – NET10	61.652	The numerical output of each of the 10 networks.
ERR1 – ERR24	00	The number of missing 5-minute data points in each of the 3-hour intervals of the delay vector (maximum of 36). This enables the user to see when missing data might be affecting the model output.
EVENTS	3	The number of networks from 10 that are predicting an event in 48-hours time.
PROTONS	-0.840	Log_{10} of the average >10MeV proton flux for the past 1-hour relative to the model run time.
Current Time	05/02/2002 1425	Same as model run time, but placed at the left hand of the table for easier reference.
XS	-8.109	Log_{10} of the average 0.4-4Å x-ray flux in W/m^2 for the past 1-hour relative to the model run time.
XL	-5.945	Log_{10} of the average 1-8Å x-ray flux in W/m^2 for the past 1-hour relative to the model run time.
Ratio	-2.168	Log_{10} of the average XS/XL ratio for the past 1-hour relative to the model run time.

Table 3 Description of information in Output.txt

The Output.txt file is tab delimited to facilitate import into other software (e.g. Microsoft Excel).

---

[*] The most recent GOES data file entry usually lags the current UTC time by around 10 minutes.

4.4 Data Quality

POPE checks GOES files for missing data and prints a message to the screen if any 5-minute values are missing from the delay vector. Each 3-hour interval is composed of 36 5-minute averages, hence there is some tolerance for loss of data. If more than 50% of the 5-minute averages are missing from an interval then a warning is printed to screen.

The precise number of missing data points in each element of the delay vector is recorded in the Output.txt file and can be used to ascertain when the network outputs may be affected by loss of data.

4.5 Maintaining the detrending function

The neural networks implemented in POPE require the XS/XL ratio to be detrended prior to being processed. Detrending is performed by a Detrender object which uses a 2nd order polynomial to model the long term solar cycle variation in the XS/XL ratio over the active years of the solar cycle (i.e. the 4 years prior to solar maximum to the 2 years after solar maximum). The coefficients that define the polynomial are contained within the Detrend.ini file. As time progresses the detrending function will fit recent data less accurately and will need to be recalculated and the coefficients updated in the Detrend.ini file.

POPE is set to use GOES-8 satellite data hence when computing the trend it is important to use x-ray data from the same satellite. To compute the trend retrieve daily x-ray averages from the SPIDR-2 data site[♠] and standardise the Julian Day number by subtracting the average Julian day and dividing by the standard deviation (this is necessary in order to calculate the coefficients accurately when fitting a function to the data). Plot XS/XL as a function of Julian Day and fit a second order polynomial to the data. (This has can be done in Microsoft Excel), which will be of form:

$$COEFF1*X^2 + COEFF2*X + CONST$$

Open the Detrend.ini file and update the information. Each coefficient will need to be altered as will the average and standard deviation of the Julian Day that was used to calculate the detrending function. There is no need to stop POPE when altering the Detrend.ini file.

In addition enter the date to which the detrending function was calculated. POPE uses this information to check to see when the detrending function runs to and issues a warning when the detrending function is more than 30 days old. Save the Detrend.ini file after altering the information. The current trending function is based on x-ray data from GOES-8 over the period 1/8/1997 to 31/12/2001.

Note that the SPIDR-2 web site only contains x-ray data up to the end of the last month. To calculate a detrending function right up to the current day it will be necessary to manually download and process 5-minnute data directly

---

[♠] http://spidr.ngdc.noaa.gov/

from the GOES ftp site[℘]. Methods within the RealTimeData object can be used to create a file of daily averages up to the current day.

In practice, the detrending function will not be significantly affected by new data within the current solar cycle (23) and experience has shown that the network outputs do not significantly alter when the detrending function is updated. The detrending function is likely to become far more critical at the beginning of a new solar cycle when there is little data to fit, and the addition of new data will significantly affect the trend.

4.6 Interpreting the predictions

Networks were trained with target values of 0 for quiet periods and 100 for events, hence a threshold value of 50 has been used to interpret the network output. For each prediction run POPE records the numerical output of each neural network and counts the number of networks that have an output of ≥50, i.e. the number of networks that are predicting an event.

By studying the output from POPE it may be found that a different interpretation of the network responses gives a better success rate, but as yet POPE has not been operating long enough to provide a useful dataset.

Testing results during development are summarised in Appendix A and show that on average each network had a 65% classification success rate, hence one would only expect the output from POPE to be correct 65% of the time at best.

POPE is intended to validate the MLP models further, and brief experience shows that 5 or more of the 10 networks regularly predict an event when none is seen, indicating a tendency to over predict. Further validation is required before any conclusions can be drawn, but realistically the models are not expected to perform any better than during testing.

4.7 Run-time problems

4.7.1 Loss of internet access

POPE does not check for an internet connection and will hang indefinitely if it runs with no internet access. If attempting to retrieve files when there is no internet access (for example, if a server is down) POPE will eventually terminate and will need to be restarted from the command line. Note that the Output.txt file is not overwritten by POPE if it is restarted.

4.7.2 Contaminated download files

It is possible for downloaded GOES data files to become corrupted during download if the connection is poor. A contaminated data file may also be downloaded if the incorrect proxy settings are used, in which case a html error page will be saved under a GOES data file name. Check the GOES data files (yyyymmdd_G8xr_5m.txt) in the 'retrieved' directory to ensure that they contain the correct data in the expected format.

[℘] http://www.sec.noaa.gov/ftpmenu/lists/xray.html

To minimise internet connection time POPE checks to see if local files from prior to 2 days ago already exist before attempting to download them. If one of these files is contaminated then it will not be replaced when POPE is restarted. To ensure that new files are downloaded clear the 'retrieved' directory before restarting POPE.

4.7.3 Incorrect proxy settings

If the PAConfig.ini file contains the wrong proxy settings then POPE will retrieve a html error page instead of the GOES data file and cause the program to crash. Check the most recent GOES data file to ensure that it contains GOES data and ensure the proxy settings are correct.

4.7.4 Incorrect System Clock

In order to retrieve current GOES data POPE obtains the current local time from the system and calculates the UTC time to generate the GOES URL filenames. If the system time is slightly fast POPE may try and retrieve a file that does not yet exist on the NGDC server if a retrieval is made shortly after UTC midnight. This will cause an error message to be saved under the GOES local filename and POPE will crash. This can be avoided by ensuring that the system clock is as accurate as possible. (Note that the system clock need only be set to local time as POPE automatically calculates the UTC time depending on the regional system settings).

5 Software Reference

This section contains an overview of each class. For a detailed explanation of the classes, methods and variables refer to the javadoc html files contained in the documentation folder of the installation. Further information is also included as comments in the source code.

5.1 Class Overviews

5.1.1 PredictionApplication.java

Contains the main method which instantiates a PredictionApplication object and starts its run thread. After execution the run thread sleeps for 60 minutes before repeating. RealTimeData objects and a PredictionApplication object are instantiated in the PredictionApplication object.

5.1.2 RealTimeData.java

Retrieves several days worth of GOES x-ray or proton real-time data relative to the current system time and can collate it into a continuous file of 5 minute averages ordering the data so that the most recent is at the top of the file. GOES URL filenames are created by taking the local system time and converting it into UTC time. A URLRetriever object is used to retrieve the files from the URLs. The class also contains methods to perform various averaging processes on GOES data files.

Constructor

`RealTimeData(int days, String datatype)`

`Days` – the number of days for which GOES data will be retrieved.

`Datatype` – determines whether x-ray or proton data files will be retrieved. Datatype must be either “x-rays” or “protons”.

Methods

`AppendOutputFile(String NewOutputLine, String Filename)`

Adds a new data line to an output file.

`CreateContinuousFile()`

Creates a continuous 5-minute file of retrieved GOES data by reading the retrieved files and writing the data to one file in reverse order so that newest data is at the top of the file.

`CreateFileStrings()`

Uses the UTC time derived from the current system time to create filenames for the remote GOES data files.

`GetUTCTime()`

Gets the local system time and generates the UTC time based on the GMT and DST offsets between the local time zone and UTC.

`lastHourOfProton()`

Calculates the latest 1-hour average of GOES proton data.

`lastHourOfXRay()`

Calculates the latest 1-hour average of X-ray Data (XS, XL and Ratio)

`log10(double x)`

Receives x and returns double value for log10(x)

`RetrieveData()`

Retrieves the GOES data files from the designated address for the required number of days and puts the data into one continuous file.

`setContinuousFile(String newname)`

Sets the name of the continuous 5-minute file produced by stitching GOES files together.

`setFileStem(String newstem)`

Sets the filestem for the data files being retrieved: e.g. _G10part_5m.txt

`setPrevDays(int days)`

Sets the number of days for which data files will be retrieved relative to the current date.

Additional Developers Methods

Several methods have been included in the class which are not used by POPE but which may be of significant use if attempting to develop or modify POPE further.

`StockDailyXrays(int startyear, int startmonth, int startday, int endyear, int endmonth, int endday)`

Calculates the daily averages of GOES X-ray data for every day between the current date and a date in the past (providing remote data files exist) This is designed to allow a data file to be built up from which a detrending function can be calculated.

`dailyXrayAverages()`

Calculates the daily average of XS, XL and XS/XL Ratio from a date two days ago relative to the current day. It was intended to run this method once every calendar day to create an updating file of x-ray data from which a detrending function could be dynamically calculated. This would prevent having to calculate the detrending function manually and alter the source code within the Detrender object.

`dailyXrayAverages(String DailyGOESFile)`

Calculates the daily x-ray average of XS and XL for the supplied data file and ammends the result to a file named DailyXrayAverages.txt. A variation of the above method.

5.1.3 URLRetriever.java

Function

GOES to a URL address and makes a local copy of the file that is found there. By default the local files are saved to a directory called 'retrieved' within the POPE installation folder.

Constructor

`URLRetriever(String URLaddress, String localcopyname)`

`URLaddress` – The URLaddress of the file to be downloaded.

`Localcopyname` – The filename which the local copy will be saved under.

Methods

`CreateDirectory()`

Method to create a new directory.

`SetProxy()`

Sets the system properties for internet connection (proxyhost and proxyport) by reading the PAconfig.ini file.

5.1.4 PredictionModel.java

Function

This class implements a prediction model of any desired window length, input average or lead-time by building the relevant delay vector from the continuous GOES x-ray data file created by the RealTimeDataObject.

The PredictionModel object contains a Detrender object to detrend the 5-minute GOES data prior to building the delay vector and NeuralNet objects which process the delay vector and perform the prediction.

Constructor

`PredictionModel(int Winlength, int AvePeriod, int ltime, String Name)`

`Winlength` - the length of the delay vector in hours that is to be built from 5-minute GOES data.

`AvePeriod` - the averaging interval in hours of the delay vector that is to be built from 5-minute GOES data

`ltime` - the lead-time of the model. Used to generate the time for which the prediction is valid.

`Name` - a reference name for the model.

Methods

`AppendResultsFile(String NewOutputLine)`

Adds a new dataline to the output file of the prediction model.

`AverageTheData()`

Computes the averages from the 5-minute data to form the required delay vector of length 'windowlength' and resolution 'AveragingPeriod'

`CreateHeader()`

Method to create a header for the output.txt file

`interpret_output(double[][] ResultArray)`

Converts the numerical network output to the text association of either an "Event" or a "Quiet" period using the value of 'threshold'.

`runModel()`

Runs the model using the latest downloaded data as a query.

`SetScales(NeuralNet NetworkArray[])`

Sets the values of M and C in y=mx+c for the scales in each NeuralNet object within a PredicitonModel.

`SetStatistics (NeuralNet NetworkArray[])`

Sets the statistics for each NeuralNet object within a PredictionModel. These are used to standardise the input data.

`SetWeights(NeuralNet NetworkArray[])`

Sets the weights for each NeuralNet object within a PredictionModel.

5.1.5 Detrender.java

Function

Will detrend an 'x' vs 'y' series using a user defined function. i.e. will calculate y - f(y) at each value of x, where f(y) is specified by the user. The class is designed to receive the 5-minute values of the XS/XL ratio, indexed with Julian Day, and detrend them.

Constructor

`Detrender(double t[], double y[], String filename)` `t[]` - the array holding the time values

`y[]` - the array holding the values f(t)

`filename` – the name of the file holding the trend function coefficients and related information.

The time values in the received array are standardised using the statistics of the detrending function. (These need to be soft coded in the source code by the user).

Methods

`SetCoefficients(double coeffone, double coefftwo, double constant)`

Sets the coefficients of the detrending function

`trend_function(double t, String name)`

Calculates the detrending function at each point in time using the coefficient file provided.

`log10(double x)`

Receives x and returns double value for $\log_{10}(x)$

5.1.6 NeuralNet.java

Function

Designed to take a delay vector and process it in the stages of: normalising, PCA projection and scaling, and then feed the processed vector as an input to an MLP neural network.

Constructor

`NeuralNet(String MName, String PCAname, boolean needPCA, int PCs)`

`MName` - the name of the model as an identifier

`PCAname` - the file holding the PCA coefficients

`needPCA` - true indicates that a PCA projection should be performed after standardizing the input vector, false indicates that no PCA projection is required.

`PCs` - the number of principal components to be used as inputs to the model.

Methods

`Denormalise_output(double normalised)`

Denormalises the output from the network to a 'real' value.

`descaler(double y)`

Receives the raw network output (double) and descales it according to the training target scale.

`Interrogate_Network(double InputVector[], double OutputVector[])`

A modified code extraction from Neuframe. This method gives the processed and scaled input vector to the network and calculates the network output.

`normalise_delay_vector(double raw_dv_vector[])`

Normalises the delay vector with using the average and standard dev of each element in the training set.

`PCAProjection(int dvlength)`

Takes a (normalised) delay vector and performs a PCA projection by reading a file of coefficients. All principal axes are projected onto. Outputs the resulting vector to a .tmp file holding the projected delay vector

`run_query(double raw_dv_vector[])`

Takes the raw input vector produced by a PredictionModel object and processes and scales the vector according to the attributes set in the constructor. The processed and scaled vector is then passed to the Interrogate_Network method.

`scaler(double x)`

Returns a scaled value based on the scale used by Neuframe ($y = mx + c$)

`setNormalStats(double ave[], double sd[])`

Sets the normalising statistics required to normalise the raw delay vector

`setScales(double m, double k)`

Sets the values for M and C in the input scaling function $y = mx + c$

`setWeights(double weights[][])`

Sets the weights matrix for the network

5.2 Disabling the repeat cycle

By default POPE runs continually in an infinite process that sleeps for sixty minutes between prediction runs. Whilst this means that POPE can be started and 'forgotten' it is usually not a good practice to run a process that has no termination, and some users may wish to disable the sleep cycle and run POPE in 'single prediction mode' via a task manager. Once the sleep-loop is removed POPE will terminate after each prediction run and will need restarting to perform another prediction.

To disable the repeating loop open the PredicitonApplication.java file in an editor and go the run() method. Delete the while(true) loop and the associated brackets and delete the 'try' and 'catch' blocks and everything contained within them.

Save the file and recompile using the command:

```
javac PredictionApplication.java
```

POPE can be started by typing

```
java PredictionApplication
```

at the MS-DOS prompt and will terminate after a prediction has been made. A task manager can be used to run the program at set times, and results will be saved to the Output.txt file. Note that the Output.txt file is not overwritten when POPE restarts.

5.3 Implementing a new prediction model

POPE has been written to allow further prediction models to be implemented with minimum effort. Provided that a model is to use GOES x-ray data as an input only new PredictionModel and new NeuralNet objects need be created. Although implementation of new models has been facilitated, modification of the existing source code will be required.

5.3.1 Building a new delay vector

The PredictionModel object will build the required delay vector from the continuous GOES 5-minute x-ray file using the delay vector characteristics that are supplied to the PredictionModel object at construction. For example, if the model being implemented requires a delay vector of length 12-hours, averaging period 2 hours and lead-time 6 hours a PredictionModel object could be instantiated using:

```
new PredictionModel(12, 2, 6, ModelName);
```

Ensure that enough GOES x-ray data is being compiled into the continuous file by the RealTimeData objects. (If not it may be necessary to change the number of days for which data is retrieved in the RealTimeData constructor). PredictionModel objects should be created in the PredictionApplication run() method.

5.3.2 Building a new MLP neural network

The NeuralNet object contains methods and variables specific to an individual MLP neural network. The MLP models implemented here were developed in Neuframe v.4 and a JAVA code extraction was performed on the trained networks.

To implement a new MLP network within POPE it is first necessary to extract the JAVA code from Neuframe. The extracted code contains errors and must be modified slightly for it to work (please refer to the comments in the source code above the Interrogate_Network method). The Neuframe extraction forms the Interrogate_Network method within NeuralNet, and the class variables from the Neuframe extraction become class variables within NeuralNet. Further comments are given in the NeuralNet.java source code.

Each NeuralNet object requires a specific set of weights, a specific set of statistics that relate to the training set and a specific set of coefficients that relate to the scaling functions. When a new NeuralNet object is instantiated these variables must be set using the relevant '`set`' methods of the NeuralNet class. The weights matrix can be obtained directly from the Neuframe extraction. The normalising statistics must be obtained from the network training set, and the scaling coefficients must be calculated from the scales used in Neuframe. (Neuframe does not display the scaling coefficients and these must be deduced manually via the information displayed in the scaling properties window. The scaling is linear and of the form $y = mx + c$ where y is the scaled output and x is the un-scaled value).

If a principal components analysis is to be performed each NeuralNet object will require a text file containing the matrix of projection coefficients. This matrix is derived from the training set. The number of principal components taken as inputs by the neural network is defined in the NeuralNet constructor. Refer to the html *javadoc* files for full details of the NeuralNet constructor.

5.3.3 Summary of how to implement a new prediction model

- Create a new PredictionModel object in the PredictionApplication `run()` method with the required delay vector characteristics.
- Within the PredictionModel constructor instantiate a new NeuralNet object.
- Set the weights, normalising statistics and scaling coefficients of the NeuralNet object.
- If a PCA projection is required ensure that a text file containing the coefficient matrix is present in the POPE directory and that it has been referenced correctly in the NeuralNet constructor.
- The model can be run using the `run_model()` method of PredictionModel but this will have to be modified slightly as it is currently written to run an array of NeuralNet objects as opposed to just one NeuralNet object.

6 Appendix A

Neural network test performance results prior to inclusion in POPE

Summary of method

A dataset was compiled consisting of 97 events and 173 quiet periods. This was randomly divided into a training set of 60 events and 60 quiet periods and a query set of 37 events and 113 quiet periods. The process was repeated 10 times to generate 10 random combinations of training and query data.

Each dimension of the training set was normalised and a PCA performed in order to project the training set onto its principal axes. The first six principal components were then taken as an input vector to a MLP network in Neuframe.

Each query set was normalised and PCA projected using the statistics and coefficients derived from the respective training set and was used to test each of the trained networks. Results are given below.

Results

Network	Events classified correctly / %	Quiet periods classified correctly / %	Overall classification Success / %
Net 1	86.5	42.5	64.5
Net 2	73.0	61.1	67.0
Net 3	81.1	45.1	63.1
Net 4	83.8	46.9	65.3
Net 5	75.7	39.8	57.7
Net 6	75.7	38.9	57.3
Net 7	81.1	55.8	68.4
Net 8	67.6	58.4	63.0
Net 9	62.2	61.9	62.1
Net 10	59.5	57.5	58.5

Reference List

1. Akasofu, S., Energy Supply Processes for Solar Flares and Magnetospheric Substorms, *Space Science Reviews*, 95, 613-621, 2001.

2. Aschwanden, M. J., M. L. Montello, B. R. Dennis, and A. O. Benz, Sequences of Correlated Hard X-Ray and Type-Iii Bursts During Solar-Flares, *Astrophysical Journal*, 440, 394-406, 1995.

3. Balch, C. and Kunches, J. SESC Methods for Proton Event Forecasts. Proceedings of Solar Terrestrial Predictions Workshop.

4. Balch, C. C., Sec Proton Prediction Model: Verification and Analysis, *Radiation Measurements*, 30, 231-250, 1999.

5. Barouch, E., M. Gros, and P. Massa, The solar longitude dependence of proton event decay, *Solar Physics*, 19, 483, 1971.

6. Bieber, J. W. and D. M. Rust, The Escape of Magnetic-Flux From the Sun, *Astrophysical Journal* , 453, 911-918, 1995.

7. Borda, R. A. F., P. D. Mininni, C. H. Mandrini, D. O. Gomez, O. H. Bauer, and M. G. Rovira, Automatic Solar Flare Detection Using Neural Network Techniques, *Solar Physics*, 206, 347-357, 2002.

8. Time Series Analysis, Forecasting and Control. Box, G. E. P. and Jenkins, G. M. Time Series Analysis, Forecasting and Control. 76. Holden-Day.

9. Burlaga, J., *Journal of Geophysical Research*, 72, 4449, 1967.

10. Cane, H. V., R. E. Mcguire, and T. T. Vonrosenvinge, 2 Classes of Solar Energetic Particle Events Associated With Impulsive and Long-Duration Soft-X-Ray Flares, *Astrophysical Journal*, 301, 448-459, 1986.

11. Cane, H. V., D. V. Reames, and T. T. Vonrosenvinge, The Role of Interplanetary Shocks in the Longitude Distribution of Solar Energetic Particles, *Journal of Geophysical Research-Space Physics*, 93, 9555-9567, 1988.

12. Cane, H. V., D. V. Reames, and T. T. Vonrosenvinge, Solar Particle Abundances at Energies of Greater Than 1-Mev Per Nucleon and the Role of Interplanetary Shocks, *Astrophysical Journal*, 373, 675-682, 1991.

13. Cane, H. V., I. G. Richardson, and T. T. Von Rosenvinge, Interplanetary Magnetic Field Periodicity of Similar to 153 Days, *Geophysical Research Letters*, 25, 4437-4440, 1998.

14. Canfield, R. C., H. S. Hudson, and A. A. Pevstov, Sigmoids as Precursors of Solar Eruptions, *IEEE Transactions on Plasma Science*, 28, 1786-1794, 2000.

15. Cargill, P. J., J. Chen, D. S. Spicer, and S. T. Zalesak, Magnetohydrodynamic Simulations of the Motion of Magnetic Flux Tubes Through a Magnetized Plasma, *Journal of Geophysical Research-Space Physics*, 101, 4855-4870, 1996.

16. Chakravorti, T. B., T. K. Das, A. K. Sen, and M. K. Dasgupta, Some Studies of Solar Proton Events in Relation to Active Region Characteristics, *Bulletin of the Astronomical Institutes of Czechoslovakia*, 42, 165-170, 1991.

17. Chen, J. and D. A. Garren, Interplanetary Magnetic Clouds - Topology and Driving Mechanism, *Geophysical Research Letters*, 20, 2319-2322, 1993.

18. Cipollini, P. C. G. D. M. G. R., Retrieval of Sea Water Optically Active Parameters from hyperspectral data by means of generalised radial basis function neural networks., *IEEE Transactions on Geoscience and Remote Sensing*, 2000.

19. Cline, K. S., N. H. Brummell, and F. Cattaneo, On the Formation of Magnetic Structures by the Combined Action of Velocity Shear and Magnetic Buoyancy, *Astrophysical Journal*, 588, 630-644, 2003.

20. Costello, K. The Costello Model. http://www.sel.noaa.gov/rpc/costello/index.html 2001.

21. Das, T. K., T. B. Chakraborty, and M. K. Dasgupta, Studies of Some Aspects of Solar Proton Events and Related Phenomena, *Bulletin of the Astronomical Institutes of Czechoslovakia*, 38, 206-210, 1987.

22. Donnelly, R. F. and L. C. Puga, 13-Day Periodicity and the Center-to-Limb Dependence of Uv, Euv, and X-Ray-Emission of Solar-Activity, *Solar Physics*, 130, 369-390, 1990.

23. Feynman, J., T. P. Armstrong, L. Daogibner, and S. Silverman, Solar Proton Events During Solar-Cycles 19, 20, and 21, *Solar Physics*, 126, 385-401, 1990a.

24. Feynman, J., T. P. Armstrong, L. Daogibner, and S. Silverman, New Interplanetary Proton Fluence Model, *Journal of Spacecraft and Rockets*, 27, 403-410, 1990b.

25. Feynman, J., G. Spitale, J. Wang, and S. Gabriel, Interplanetary Proton Fluence Model - Jpl 1991, *Journal of Geophysical Research-Space Physics*, 98, 13281-13294, 1993.

26. Gabriel, S., R. Evans, and J. Feynman, Periodicities in the Occurrence Rate of Solar Proton Events, *Solar Physics*, 128, 415-422, 1990.

27. Gabriel, S. B. and J. Feynman, Power-Law Distribution for Solar Energetic Proton Events, *Solar Physics*, 165, 337-346, 1996.

28. Gabriel, S. B., Feynman, J., and Spitale G. Solar Energetic Particle Events: Statistical Modelling and Prediction. ESA Symposium Proceedings on 'Environment Modelling for Space-based Applications. SP-392. 96.

29. Garcia, H. A., Temperature and Hard X-Ray Signatures for Energetic Proton Events, *Astrophysical Journal*, 420, 422-432, 1994a.

30. Garcia, H. A., Temperature and Emission Measure From Goes Soft-X-Ray Measurements, *Solar Physics*, 154, 275-308, 1994b.

31. Garcia, H. A., Farnik, F., and Kiplinger, A. L. Hard X-Ray Spectroscopy for Proton Flare Prediction Spie - Int Soc Optical Engineering. Proceedings of the Society of Photo-Optical Instrumentation Engineers (Spie).

32. Garcia, H. A. and Kiplinger, A. L. Low-Temperature Soft X-Ray Flares, Spectrally Hardening Hard X- Ray Flares, and Energetic Interplanetary Protons. 95, 91-99. 96. Astronomical Society of the Pacific Conference Series.

33. Gleisner, H. and H. Lundstedt, A Neural Network-Based Local Model for Prediction of Geomagnetic Disturbances, *Journal of Geophysical Research-Space Physics*, 106, 8425-8433, 2001.

34. Gothoskar, P. and S. Khobragade, Detection of Interplanetary Activity Using Artificial Neural Networks, *Monthly Notices of the Royal Astronomical Society*, 277, 1274-1278, 1995.

35. Haykin, S. Neural Networks, a comprehensive foundation. Haykin, S. Neural Networks, a comprehensive foundation. 99. Prentice Hall.

36. Heckman, G. Prediction of Solar Particle Events for Exploration Class Missions. Biological Effects and Physics of Solar and Galactic Cosmic Radiation. 243, 89-100. 93. Plenum Press Div Plenum Publishing Corp. Nato Advanced Science Institutes Series, Series a, Life Sciences.

37. Heckman, G. R. Solar Proton Event Forecasts. Interplanetary Particle Environment Conference Proceedings . pp 91-100. 88. JPL publication 88-28.

38. Heckman, G. R., J. M. Kunches, and J. H. Allen, Prediction and Evaluation of Solar Particle Events Based on Precursor Information, *Advances in Space Research.*, Vol. 12, (2)313-(2)320, 1991.

39. Heckman, G. R., W. J. Wagner, J. W. Hirman, and J. M. Hunches, Strategies for Dealing with Solar Particle Events in Missions Beyond the Magentosphere, *Advances in Space Research*, 9, 275-280.

40. Horne R. B. Space Weather Parameters Required by the Users and Synthesis of User Requirements wp1300 and wp1400 (Version 3.1). 2001.

41. Huang, Y. N. and S. J. Wang, Two-Step Coronal Transport of Solar Flare Particles From Magnetic Multipolarity Sources in a Flare Region, *Chinese Physics Letters*, 18, 304-306, 2001.

42. Kahler, S. W., Radio-Burst Characteristics of Solar Proton Flares, *Astrophysical Journal*, 261, 710-719, 1982a.

43. Kahler, S. W., The Role of the Big Flare Syndrome in Correlations of Solar Energetic Proton Fluxes and Associated Microwave Burst Parameters, *Journal of Geophysical Research-Space Physics*, 87, 3439-3448, 1982b.

44. Kahler, S. W., Solar-Flares and Coronal Mass Ejections, *Annual Review of Astronomy and Astrophysics*, 30, 113-141, 1992.

45. Kahler, S. W., N. R. Sheeley, R. A. Howard, M. J. Koomen, D. J. Michels, R. E. Mcguire, T. T. Vonrosenvinge, and D. V. Reames, Associations Between Coronal Mass Ejections and Solar Energetic Proton Events, *Journal of Geophysical Research-Space Physics* , 89, 9683-9693, 1984.

46. Kallenrode, M. B. Particle Acceleration at Interplanetary Shocks - Observations at a Few Tens of Kev Vs Some Tens of Mev. 15(8-9), 375-384. 95. Advances in Space Research.

47. Kallenrode, M. B., G. Wibberenz, H. Kunow, R. Mullermellin, V. Stolpovskii, and N. Kontor, Multi-Spacecraft Observations of Particle Events and Interplanetary Shocks During November December 1982, *Solar Physics*, 147, 377-410, 1993.

48. Kalman, B., Submergence of Magnetic Flux in Interaction of Sunspot Groups, *Astronomy & Astrophysics*, 371, 731-737, 2001.

49. Kane, R. P., E. R. Depaula, and N. B. Trivedi, Variations of Solar Euv, Uv and Ionospheric Fof2 Related to the Solar Rotation Period, *Annales Geophysicae-Atmospheres Hydrospheres and Space Sciences*, 13, 717-723, 1995.

50. King, J. H., Solar Proton Fluences for 1977-1983 space missions, *Journal of Spacecraft and Rockets*, 11, 401, 1974.

51. Kiplinger, A. L., Comparative-Studies of Hard X-Ray Spectral Evolution in Solar- Flares With High-Energy Proton Events Observed at Earth, *Astrophysical Journal*, 453, 973-986, 1995.

52. Kunches, J. M. and R. D. Zwickl, The Effects of Coronal Holes on the Propagation of Solar Energetic Protons, *Radiation Measurements*, 30, 281-286, 1999.

53. Lundstedt, H., H. Gleisner, and P. Wintoft, Operational Forecasts of the Geomagnetic Dst Index, *Geophysical Research Letters*, 29, art. no.-2181, 2002.

54. Meyer, J., G. Wibberenz, and M. B. Kallenrode, Time-Development of Proton Energy-Spectra in Solar Energetic Particle Events, *Advances in Space Research*, 13, 363-366, 1993.

55. Miller, J. A., Particle Acceleration in Impulsive Solar Flares, *Space Science Reviews*, 86, 79-105, 1998.

56. Montgomery D.C., Johnson, L. A., and Gardiner J.S. Forecasting and Time Series Analysis. Montgomery D.C., Johnson, L. A., and Gardiner J.S. Forecasting and Time Series Analysis. 90. McGraw-Hill.

57. Moore, R. L., Evidence That Magnetic Energy Shedding in Solar Filament Eruptions Is the Drive in Accompanying Flares and Coronal Mass Ejections, *Astrophysical Journal*, 324, 1132-1137, 1988.

58. Mursula, K. and B. Zieger, The 13.5-Day Periodicity in the Sun, Solar Wind, and Geomagnetic Activity: the Last Three Solar Cycles, *Journal of Geophysical Research-Space Physics*, 101, 27077-27090, 1996.

59. O'sullivan, D. and D. Zhou, Overview and Present Status of the European Commission Research Programme, *Radiation Protection Dosimetry*, 86, 279-283, 1999.

60. O'sullivan, D., D. Zhou, and E. Flood, Investigation of Cosmic Rays and Their Secondaries at Aircraft Altitudes, *Radiation Measurements*, 34, 277-280, 2001.

61. O'sullivan, D., D. Zhou, W. Heinrich, S. Roesler, J. Donnelly, R. Keegan, E. Flood, and L. Tommasino, Cosmic Rays and Dosimetry at Aviation Altitudes, *Radiation Measurements*, 31, 579-584, 1999.

62. Oliver, R., J. L. Ballester, and F. Baudin, Emergence of Magnetic Flux on the Sun as the Cause of a 158-Day Periodicity in Sunspot Areas, *Nature*, 394, 552-553, 1998.

63. Paola, J. D. S. R. A., A review and analysis of backpropagation neural networks for classification of remotely-sensed multi-spectral imagery, *International Journal of Remote Sensing*, 16, 3033-3058, 1995.

64. Priest, E. R. and T. G. Forbes, The Magnetic Nature of Solar Flares, *Astronomy and Astrophysics Review*, 10, 313-377, 2002.

65. Reames, D. V., Bimodal Abundances in the Energetic Particles of Solar and Interplanetary Origin, *Astrophysical Journal*, 330, L71-L75, 1988.

66. Reames, D. V., Acceleration of Energetic Particles by Shock-Waves From Large Solar-Flares, *Astrophysical Journal*, 358, L63-L67, 1990.

67. Reames, D. V., Solar Energetic Particles - a Paradigm Shift, *Reviews of Geophysics*, 33, 585-589, 1995.

68. Reames, D. V., Particle Acceleration at the Sun and in the Heliosphere, *Space Science Reviews*, 90, 413-491, 1999.

69. Reames, D. V., Magnetic Topology of Impulsive and Gradual Solar Energetic Particle Events, *Astrophysical Journal*, 571, L63-L66, 2002.

70. Reames, D. V., L. M. Barbier, T. T. Vonrosenvinge, G. M. Mason, J. E. Mazur, and J. R. Dwyer, Energy Spectra of Ions Accelerated in Impulsive and Gradual Solar Events, *Astrophysical Journal*, 483, 515-522, 1997.

71. Rosenblatt, F., The Perceptron: A probabilistic model for information storage and organisation in the brain. *Psychological Review*, 65, 306-408, 1958.

72. Shea, M. A. and D. F. Smart, A Summary of Major Solar Proton Events, *Solar Physics*, 127, 297-320, 1990.

73. Shea, M. A. and D. F. Smart, Recent and Historical Solar Proton Events, *Radiocarbon*, 34, 255-262, 1992.

74. Shea, M. A. and Smart, D. F. Significant Proton Events of Solar-Cycle-22 and a Comparison With Events of Previous Solar-Cycles. *Advances in Space Research* 14(10), 631-638. 94. .

75. Shea, M. A. and Smart, D. F. Solar Proton Fluxes as a Function of the Observation Location With Respect to the Parent Solar-Activity. *Advances in Space Research* 17(4/5), 225-228. 95a.

76. Shea, M. A. and D. F. Smart, History of Solar Proton Event Observations, *Nuclear Physics B*, 16-25, 1995b.

77. Shea, M. A. and Smart, D. F. Space Weather: the Effects on Operations in Space. 22(1), 29-38. 98. *Advances in Space Research.*

78. Shea, M. A., D. F. Smart, J. H. Allen, and D. C. Wilkinson, Spacecraft Problems in Association With Episodes of Intense Solar-Activity and Related Terrestrial Phenomena During March 1991, *Ieee Transactions on Nuclear Science*, 39, 1754-1760, 1992.

79. Sheeley, N. R., R. A. Howard, M. J. Koomen, and D. J. Michels, Associations Between Coronal Mass Ejections and Soft-X-Ray Events, *Astrophysical Journal*, 272, 349-354, 1983.

80. Smart, D. F. and M. A. Shea, Solar Proton Events During the Past 3 Solar-Cycles, *Journal of Spacecraft and Rockets*, 26, 403-415, 1989.

81. Smith, Z., M. Dryer, E. Ort, and W. Murtagh, Performance of interplanetary shock prediction models: STOA and ISPM, *Journal of Atmospheric and Solar-Terrestrial Physics*, 1265-1274, 2000.

82. Stassinopoulos, E. G., G. J. Brucker, J. N. Adolphsen, and J. Barth, Radiation-Induced Anomalies in Satellites, *Journal of Spacecraft and Rockets*, 33, 877-882, 1996.

83. Stassinopoulos, E. G., G. J. Brucker, D. W. Nakamura, C. A. Stauffer, G. B. Gee, and J. L. Barth, Solar Flare Proton Evaluation at Geostationary Orbits for Engineering Applications, *IEEE Transactions on Nuclear Science*, 43, 369-382, 1996.

84. Stewart, *Phil. Trans. Royal Soc. London*, 407, 1861.

85. Tanaka, K., Impact of X-Ray Observations From the Hinotori Satellite on Solar-Flare Research, *Publications of the Astronomical Society of Japan*, 39, 1-45, 1987.

86. Tranquille, C., Solar Proton Events and Their Effect on Space Systems, *Radiation Physics and Chemistry*, 43, 35-45, 1994.

87. Uddin, W., M. C. Pande, and V. K. Verma, On the Association of Solar Proton Events With Other Flare Manifestations, *Bulletin of the Astronomical Institutes of Czechoslovakia*, 41, 374-380, 1990.

88. Vampola, A. L., M. Lauriente, D. C. Wilkinson, J. Allen, and F. Albin, Single Event Upsets Correlated With Environment, *Ieee Transactions on Nuclear Science*, 41, 2383-2388, 1994.

89. Vrsnak, B. and S. Lulic, Formation of Coronal Mhd Shock Waves - I. The Basic Mechanism, *Solar Physics*, 196, 157-180, 2000a.

90. Vrsnak, B. and S. Lulic, Formation of Coronal Mhd Shock Waves - Ii. The Pressure Pulse Mechanism, *Solar Physics*, 196, 181-197, 2000b.

91. Wagner, W. J., Observations of 1-8A Solar X-ray Variability During Solar Cycle 21, *Advances in Space Research*, 8, (7)67-(7)76, 1988.

92. Watanabe, S., E. Sagawa, K. Ohtaka, and H. Shimazu, Prediction of the Dst Index From Solar Wind Parameters by a Neural Network Method, *Earth Planets and Space*, 54, 1263-1275, 2002.

93. Widrow, B. and M. E. Hoff, Adaptive Switching Circuits, *IRE WESCON Convention record*, 96-104, 1960.

94. Wilkinson, D. C., National Oceanic and Atmospheric Administrations Spacecraft Anomaly Data-Base and Examples of Solar-Activity Affecting Spacecraft, *Journal of Spacecraft and Rockets*, 31, 160-165, 1994.

95. Wilkinson, D. C., S. C. Daughtridge, J. L. Stone, H. H. Sauer, and P. Darling, Tdrs-1 Single Event Upsets and the Effect of the Space Environment, *Ieee Transactions on Nuclear Science*, 38, 1708-1717, 1991.

96. Wu, G. P., D. Y. Wang, A. A. Xu, and Y. H. Tang, A Study of the Mechanism of Acceleration of He-3 and Heavy Ions by Alfven Turbulence in Impulsive Flares, *Chinese Astronomy and Astrophysics*, 26, 301-307, 2002.

97. Xapsos, M. A., G. P. Summers, and E. A. Burke, Probability Model for Peak Fluxes of Solar Proton Events, *Ieee Transactions on Nuclear Science*, 45, 2948-2953, 1998.

98. Xapsos, M. A., G. P. Summers, P. Shapiro, and E. A. Burke, New Techniques for Predicting Solar Proton Fluences for Radiation Effects Applications, *Ieee Transactions on Nuclear Science*, 43, 2772-2777, 1996.

Bibliogrpahy

Applied Forecasting Methods
Thomopoulus, N.T.
P Prentice-Hall Inc.
ISBN 0-13-040139

Statistical Methods for Engineers
McHuen, R.H.
p. Prentice-Hall Inc.
ISBN 0-13-844903-1

Neural Networks 2nd Edition, A Comprehensive foundation
Haykin, S.
Prentice Hall
ISBN 0-13-273350-1

The Essence of Neural Networks
Callan, R.
Prentice Hall
ISBN 0-13-908732-X

Forecasting and Time Series Analysis 2nd Edition
Montgomery, D.C., Johnson, L.A., Gardiner, J.S.
McGraw-Hill
ISBN 0-07-042858-1

Practical Time Series
Janacek, G.
Arnold
ISBN 0-340-71999-0

Beginners Guide to the Sun
Taylor, P., Hendrickson, N.L.
Kalmback Books
ISBN 0-913135-23-2

Guide to The Sun
Phillps, K.J.H.
Cambridge University Press
ISBN 0-521-39788-X

Space Mission Analysis and Design 2nd Edition
Larson, W.J., Wertz, J.R.
Kluwer Academic Publishers
ISBN 1-881883-01-9

The Behaviour of Systems in the Space Environment
Editors: Dewitt, RN, Duston, D., Hyder, A.K.
NATO ASI Series
Kluwer Academic Publishers
ISBN 0-7923-245