

## RESEARCH ARTICLE

10.1029/2018GC007714

## Key Points:

- MagPySV is an open-source Python package for creating reproducible high-resolution time series of internal secular variation
- Implemented denoising method uses principal component analysis to characterize external contamination in different geographic regions
- Denoised data from MagPySV and their application to geomagnetic jerks presented in case studies for Europe and high northern latitudes

## Correspondence to:

G. A. Cox,  
gracecox@cp.dias.ie

## Citation:

Cox, G. A., Brown, W. J., Billingham, L., & Holme, R. (2018). MagPySV: A Python package for processing and denoising geomagnetic observatory data. *Geochemistry, Geophysics, Geosystems*, 19, 3347–3363. <https://doi.org/10.1029/2018GC007714>

Received 29 MAY 2018

Accepted 27 AUG 2018

Accepted article online 31 AUG 2018

Published online 22 SEP 2018

©2018. The Authors.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

## MagPySV: A Python Package for Processing and Denoising Geomagnetic Observatory Data

G. A. Cox<sup>1,2</sup>, W. J. Brown<sup>3</sup>, L. Billingham<sup>3,4</sup>, and R. Holme<sup>1</sup>

<sup>1</sup>Department of Earth, Ocean and Ecological Sciences, University of Liverpool, Liverpool, UK, <sup>2</sup>Geophysics Section, Dublin Institute for Advanced Studies, Dublin, Ireland, <sup>3</sup>British Geological Survey, Lyell Centre, Edinburgh, UK, <sup>4</sup>TravelNest Limited, Edinburgh, UK

**Abstract** Measurements obtained at ground-based observatories are crucial to understanding the geomagnetic field and its secular variation (SV). However, current data processing methods rely on piecemeal closed-source codes or are performed on an ad hoc basis, hampering efforts to reproduce data sets underlying published results. We present MagPySV, an open-source Python package designed to provide a consistent and automated means of generating high-resolution SV data sets from hourly means distributed by the Edinburgh World Data Centre. It applies corrections for documented baseline changes, and optionally, data may be excluded using the *ap* index, which removes effects from documented high solar activity periods such as geomagnetic storms. Robust statistics are used to identify and remove outliers. Developing existing denoising methods, we use principal component analysis of the covariance matrix of residuals between observed SV and that predicted by a global field model to remove a proxy for external field contamination from observations. This method creates a single covariance matrix for all observatories of interest combined and applies the denoising to all locations simultaneously, resulting in cleaner time series of the internally generated SV. In our case studies, we present cleaned data in two geographic regions: monthly first differences are used to investigate geomagnetic jerk morphology in Europe, an area previously well-studied at lower resolution, and annual differences are investigated for northern high latitude regions, which are often neglected due to their high noise content. MagPySV may be run on the command line or within an interactive Jupyter notebook; two notebooks reproducing the case studies are supplied.

## 1. Introduction

These external fields also act to induce secondary fields in the Earth, which vary in time with the inducing field. A major challenge in using geomagnetic observations is the difficulty in separating variations in the core-generated field from those of external sources at overlapping time intervals, which often obscure fine-scale details. Observatory annual means (OAMs) are widely used to study time changes of the geomagnetic field (secular variation [SV]) and their underlying physical processes in the outer core, in which motions of molten iron generate the field through dynamo action. However, their temporal resolution is not sufficient for studies of rapid changes in the core-generated SV, such as geomagnetic jerks (see Mandaia et al., 2010, for a review of this topic). Observatory monthly means are better suited for studying rapid core dynamics but are highly sensitive to measurement errors due to their high temporal resolution. Both means suffer from significant external magnetic field contamination (*noise, in the context of internal field studies*). The observed geomagnetic field is a combination of various sources, both internal and external to the Earth. The dynamo-generated field (time-varying) and the crustal field arising from local geology (assumed steady on observation time scales) are the main internal field sources. Examples of external magnetic field sources include electric current systems in the ionosphere and magnetosphere and their interactions with the solar wind. In this paper, we present MagPySV, an open-source Python package designed to compile, process, and denoise observatory data, resulting high-resolution SV series with documented baseline changes outliers and external contamination removed. A method for denoising observatory means using principal component analysis (PCA), based on work by Wardinski and Holme (2011), Brown et al. (2013), and Feng et al. (2018), is presented in this work and implemented in the software. This denoising method permits the use of higher temporal resolution data than have previously been used and allows better characterization of noise in different geographic regions than prior studies. We illustrate this point with two case studies in this paper: one for European observatories and a second for northern high latitude observatories. Two example Jupyter note-

books are provided with MagPySV, which allows the reader to reproduce the analyses and figures presented in this paper.

### 1.1. Data Holdings

The primary data holding of geomagnetic observatory measurements is the Geomagnetic Data Master Catalogue of the World Data Centre (WDC) for Geomagnetism in Edinburgh (hereafter WDCE), a regular member of the International Council for Science World Data System. This is in essence a repository to ensure the long-term availability of data collected at observatories (since around 1840) and is not a processing facility, although some quality control (QC) is performed to correct or report anomalies. Data are reported by standard as 1-min and 1-hr means.

The International Real-Time Magnetic Observatory Network (INTERMAGNET) is a body that promotes standards in geomagnetic observatory operation including QC and data delivery time, with a focus on real-time collection of data through five geographically distributed Geomagnetic Information Nodes. Since 2013, INTERMAGNET classifies data into four categories according to their quality (listed from the lowest to highest): *variation*, *provisional*, *quasi-definitive* (QD), and *definitive*, at 1-min and 1-s resolution. Of these, QD and definitive data are most suitable for research purposes, with QD data being delivered within 3 months of acquisition, corrected to provisional baselines, and very close to the final definitive values determined on a longer time scale (Clarke et al., 2013; Peltier & Chulliat, 2010). Data from observatories that meet the standards set by INTERMAGNET are hosted for public online access.

All definitive data from INTERMAGNET observatories are deposited in WDCE, as well as definitive data from unaffiliated observatories. Various subsets of data can also be accessed through WDCs Kyoto, Mumbai, and Moscow.

Further to these *raw* data holdings, derived products are produced and distributed by various institutions. The most relevant here are the OAM values at WDCE distributed by British Geological Survey (BGS; [http://www.geomag.bgs.ac.uk/data\\_service/data/annual\\_means.shtml](http://www.geomag.bgs.ac.uk/data_service/data/annual_means.shtml)), the monthly mean values produced by Bureau Central de Magnétisme Terrestre (BCMT; Chulliat & Telali, 2007; <ftp://ftp.bcmt.fr/>), and the *AUX\_OBS* data sets produced by BGS to aid the European Space Agency Swarm mission ([ftp://ftp.nerc-murchison.ac.uk/geomag/Swarm/AUX\\_OBS/](ftp://ftp.nerc-murchison.ac.uk/geomag/Swarm/AUX_OBS/)).

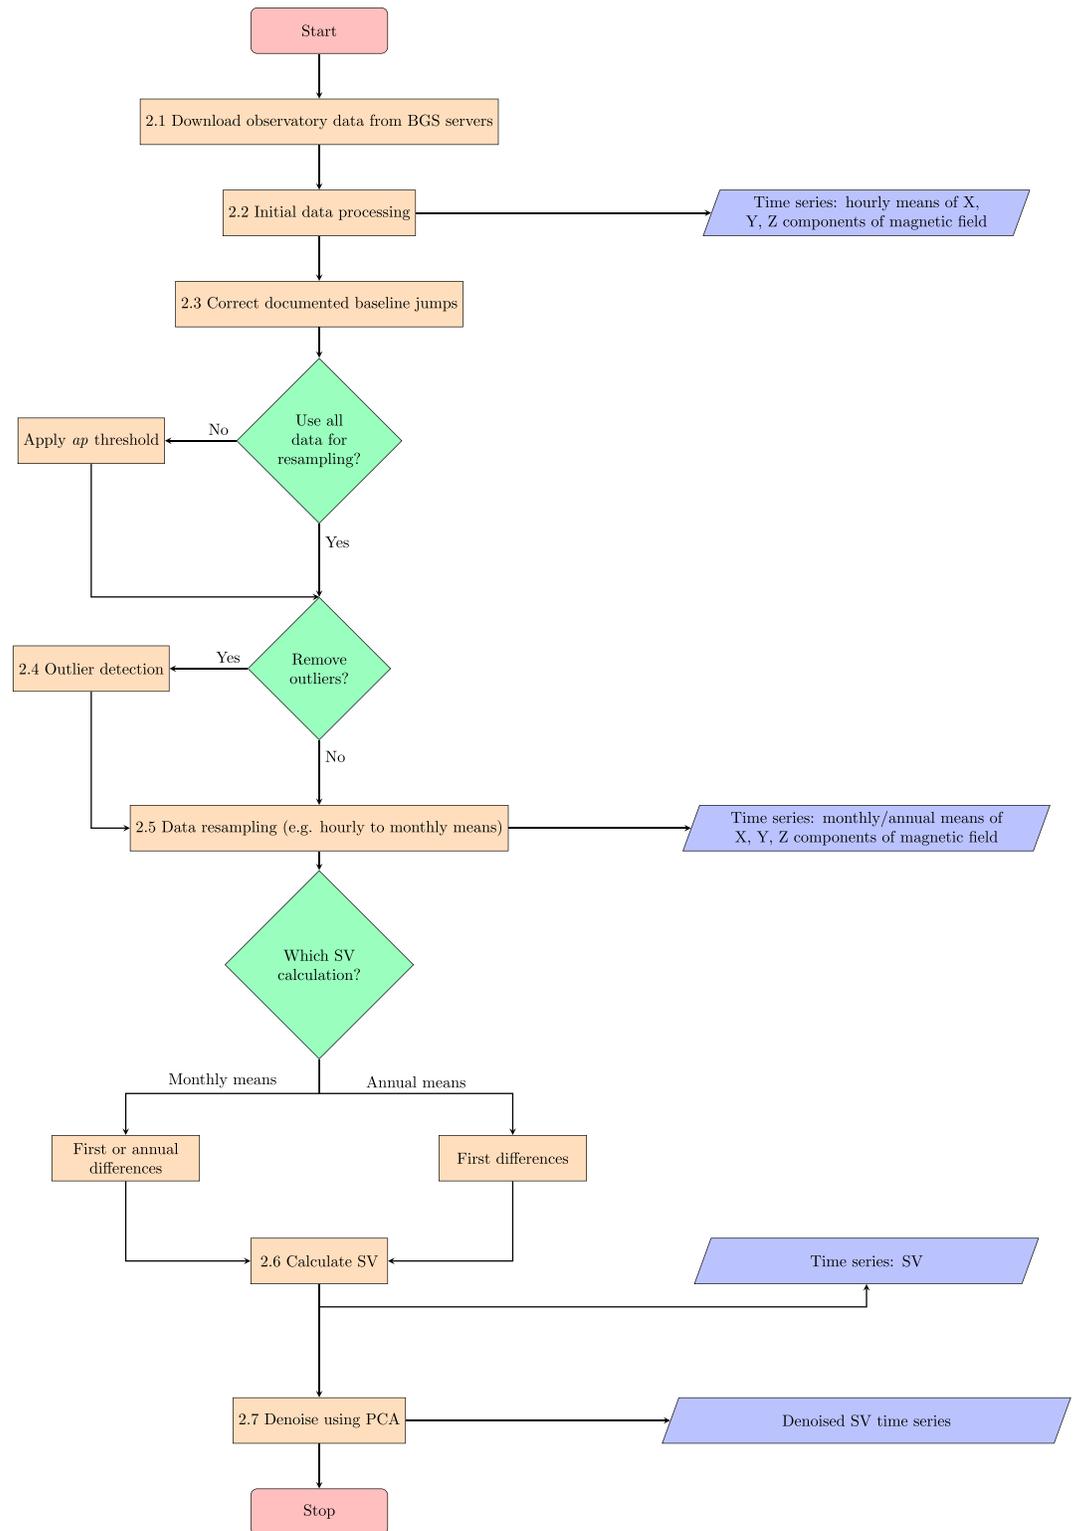
The OAM files also contain information regarding all documented baseline discontinuities at each observatory; the discontinuity values are also given on the OAM pages of the BGS website. The BCMT supplies annual-, monthly-, daily-, hourly-, and minute-mean and second cadence data for the BCMT observatory network. The monthly-mean database is derived for all INTERMAGNET observatories from WDC hourly-mean holdings with a QC and calculation procedure described by Chulliat and Telali (2007). No denoising methods to remove external noise contamination were implemented in that work, other than averaging the raw data over 1 month.

The *AUX\_OBS* hourly data set (called *AUX\_OBS\_2*) originally covered the era of near-continuous satellite observation from 1997 to present, updated every 3 months, to aid global field modeling efforts. Both QD and definitive standard data are collated on a three monthly basis for as many observatories as possible and a thorough QC procedure applied, as described by Macmillan and Olsen (2013). This procedure includes the application of all known baseline discontinuities and the splitting of records where unknown discontinuities are identified. Since 2017, the hourly *AUX\_OBS\_2* set has been extended back to 1957 and augmented by daily updates minute mean *AUX\_OBSM2* (from 1997) and 1-s *AUX\_OBS2* (from 2013) sets. Note that in contrast to the WDCE data, which are given in the geodetic frame, all *AUX\_OBS* data are given in the spherical frame. Note also that they are given in a different file format.

Initially, we here choose to use the definitive hourly mean data, held at WDCE, in order to access the widest ranging historical data set at a cadence suited to internal field studies and the largest geographic spread of observatories. Future updates will look to extend access to other data holdings and to implement access to other cadence data.

## 2. Method

In order to go from raw magnetic field data to denoised SV series suitable for studying rapid observed features and core dynamics, several processing steps are required. Figure 1 shows the workflow implemented by MagPySV; each step is described in the following text.



**Figure 1.** Flowchart showing the basic workflow for processing and denoising raw World Data Centre for Geomagnetism in Edinburgh geomagnetic hourly data. Orange boxes represent processes in the workflow, green diamonds represent decisions made by the user, and blue trapezoids indicate the output from the code at various stages of the workflow. The numbers in the process boxes correspond to the sections of this paper in which the processes are described. BGS = British Geological Survey; SV = secular variation; PCA = principal component analysis.

### 2.1. Obtaining WDC Data

Geomagnetic data hosted by WDCE are accessed through a utility of the BGS Geomagnetism Data Portal (GDP) to the WDC master catalog (<http://wdc.bgs.ac.uk/dataportal/>). The GDP provides an interactive web application, with a graphical user interface for searching, visualization, and downloads, as well as programmatic access to the supporting web service via RESTful API. This allows a user to directly send an HTTP “get” request to retrieve an HTML form and then an HTTP “post” request to populate the form with the details of data files requested. This form entry results in a returned ZIP archive containing the data set requested, which is then unpacked and the individual files written to disk for the user.

In order to use this functionality of the web service, MagPySV depends on a Python library developed in tandem at BGS, which is installed as part of the MagPySV installation. This library provides the ability to request a data download, through a single Python function call, of minute or hour cadence data for a user-specified list of observatories and a user-specified date range. Further details of this code are given in section 5.

### 2.2. Initial Data Processing

All raw hourly data at WDCE are stored in a specific file format called *WDC format*, a detailed description of which may be found at <http://www.wdc.bgs.ac.uk/catalog/format.html>. Each file contains data for a single observatory in a single year, and all field components are given within a single file. The first average value in the file represents half past midnight in coordinated universal time (00:30:00 UTC) of the specified day, and the final hourly value is for 23:30:00 UTC of the day. The final value in the file, representing the daily average, is neglected as these are not calculated by the BGS/WDC but by individual observatories, and some values may be unreliable (S. Macmillan, personal communication, May 29, 2017). Magnetic field values (declination/inclination values) are converted to nanoteslas (nT; degrees) using the tabular base and values. Magnetic declination ( $D$ ) and the horizontal field component ( $H$ ) are converted into the geodetic  $X$  (north) and  $Y$  (east) field components using the relations

$$X = H \cos D, Y = H \sin D. \quad (1)$$

The raw data files are processed and placed into a Pandas dataframe (McKinney, 2010) indexed by datetime objects representing each hourly average and containing columns for the  $X$ ,  $Y$ , and  $Z$  magnetic field components. Missing values, indicated as 9999 in the WDC files, are represented by Not a Number (NaN) values in the dataframe.

### 2.3. Correction of Baseline Discontinuities

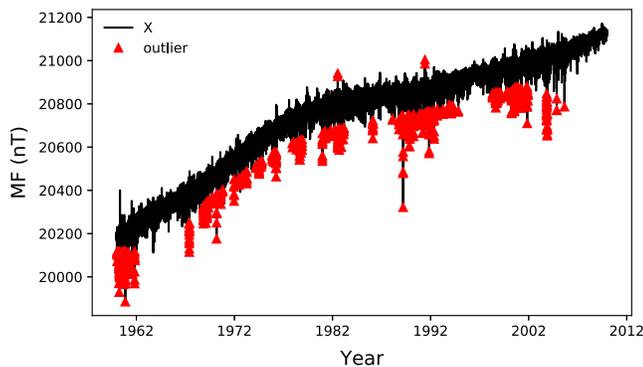
Discontinuities in an observatory baseline may appear in the data due to, for example, changes to instrumentation or moving an instrument or site. The BGS holds a list of all reported baseline discontinuities for each observatory; these are included within the OAM pages on the BGS website. We have collated these documented baseline changes up to the current day into a single file containing the observatory name, event year, and the  $X$ ,  $Y$ , and  $Z$  discontinuity values and included the current version in the Python package. The correction is applied by subtracting the given values from all preceding data. The WDC relies upon individual observatories to QC the data, keeps detailed records of any disturbance to the baseline, and communicates their dates and magnitudes. The quality and quantity of the such information vary between different observatories and through time, and there remain undocumented discontinuities in the data. Where such undocumented features are identified in the data, their times and magnitudes can be added to the file for automated correction.

### 2.4. Outlier Detection

Outliers (*spikes*) remain in the data for several reasons, such as transient external activity that is not captured by a magnetic field index, instrument errors, or the temporary presence of magnetic material near an instrument. The latter two cases should be identified and removed from the data as they do not relate to any physical processes of the magnetic field, but all are considered noise when the internal field is the primary interest. Identifying outliers in time series data is a large ongoing topic of research in many different fields, and various statistical methods have been proposed, see Barnett and Lewis (1974) and Hodge and Austin (2004) for comprehensive reviews of this topic. One method that has been applied in various contexts is the median absolute deviation (MAD) from the median, defined as

$$\text{MAD} = \text{median}(|x_i - \text{median}(x)|), \quad (2)$$

for a set of observations  $x_i$  (Hampel, 1974). The MAD is a measure of statistical dispersion and is insensitive



**Figure 2.** Hourly X component of magnetic field (MF) at Chambon-la-Forêt, France (black line), with identified outliers (red triangles).

to the presence of outliers, making it more robust than the widely used *three sigma rule*, which consists of determining an interval spanning over the mean plus or minus three standard deviations, since the mean and standard deviation are themselves sensitive to outliers. Using the MAD, we determine the rejection criterion of a value  $x_i$  as

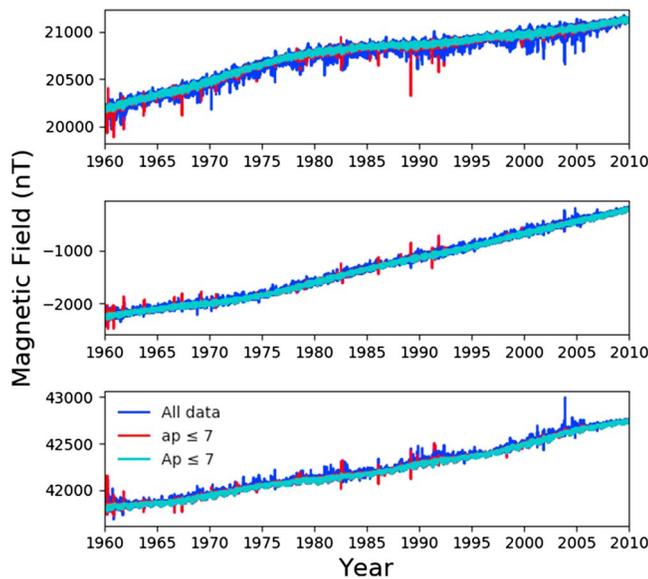
$$\frac{|x_i - \text{median}(x)|}{\text{MAD}} > \alpha, \quad (3)$$

where  $\alpha$  is a user-specified threshold. Geomagnetic time series are often long (several decades) and highly variable within that period. These variations imply that the median naturally changes over the whole series length, so we do not think it appropriate to use a single value of this quantity in the above criterion as is typical. We instead use running median calculations, with a window length specified by the user. Identified outliers are set to NaN in the dataframe. There is an option to plot the identified outliers in the time series; Figure 2 shows an example plot of hourly means

for Chambon-la-Forêt (CLF) observatory with a window length of 10 years and a threshold of 10. Given that many of the identified outliers are negative, it is likely that these are caused by strong geomagnetic storms as opposed to data errors. Note that in this paper, we define *outliers* in the statistical sense according to the MAD criterion, regardless of whether they are due to measurement error or external magnetic field processes. We recommend removing outliers from hourly field data before resampling to monthly means so as to avoid biasing entire months of data, though the user may perform outlier removal on field/SV data of any frequency.

### 2.5. Data Resampling

At this stage, the dataframes contain hourly averages of magnetic field observations in the X, Y, and Z directions indexed by datetime stamps. Data at this frequency are not generally used for studies of the outer core and internally generated SV, rather monthly or annual means, although see, for example, the GRIMM (Lesur et al., 2015) or BGS model series (Hamilton et al., 2015). The datetime objects allow easy data resampling of the hourly values at a user-specified frequency. For annual means, the datetime objects of resampled data are set to the first day of July, the midpoint of the year, while for monthly means, the datetime objects are set to the 15th day of the month. Optionally, the user may exclude data from the averaging calculation based on values of a planetary magnetic field index. Various magnetic indices are used to identify “quiet” and “disturbed” periods of data, each based on different measurements and related to the activity of different processes in the ionosphere and magnetosphere. The quasi-logarithmic 3-hourly  $K_p$  index is commonly used for data rejection (see, e.g., Rangarajan, 1989), providing an indicator for overall global geomagnetic activity from external sources, and is based on average values of the irregular disturbance levels in the horizontal field components at 13 selected midlatitude magnetic observatories. Directly related to, and derived from, the  $K_p$  index is the linear  $ap$  index, which is also provided on a 3-hourly basis. Low  $ap$  values indicate low disturbance levels, while the highest values indicate geomagnetic storms. Excluding data for which  $ap$  is high removes the effects of documented storms and other documented signals in external fields. Note, however, that the  $K_p$  network is heavily weighted toward Europe and Northern America, and using this index for data selection at magnetic observatories outside these two regions is often inappropriate. An often used criterion for geomagnetic field modeling, which uses only quiet data, is  $ap \leq 7$  (corresponding to  $K_p \leq 2^\circ$ ), which is satisfied for approximately 50% of all data, though this percentage is lower during years of solar maximum and higher during years of solar minimum (Hulot et al., 2007). MagPySV uses the definitive  $ap$  index provided by GFZ, Potsdam, at <ftp://ftp.gfz-potsdam.de/pub/home/obs/kp-ap/>. The index is used to match datetimes and drop magnetic field data from the dataframe at times where the documented  $ap$  is above the user-specified threshold. We have included two different files containing  $ap$  data with the example Jupyter notebooks: one in which the 3-hourly values are repeated so that each hour has its own value and the threshold is applied on an hourly basis and another in which the daily average of the 3-hourly values (i.e., the  $Ap$  index) is set to each hourly value so that entire days of data are dropped when the threshold is applied. Figure 3 shows the hourly magnetic field for all data (blue),  $Ap \leq 7$  (cyan), and  $ap \leq 7$  (red) at Chambon-la-Forêt, France, which shows fewer spikes and a general reduction in noise with an  $ap$  threshold applied. The largest spikes in the figure correspond to documented geomagnetic storms; however, not all of these are captured by the index. For example, the March 1989 storm, best seen in the X component, is removed with the applied  $Ap$  threshold (entire days removed) but not the applied  $ap$  threshold (only certain hours of the day removed). Since many of the surrounding data are removed, such large spikes bias the monthly mean and result in large SV outliers.



**Figure 3.** Hourly magnetic field at Chambon-la-Forêt, France. The top panel is the X component, middle panel is Y, and bottom panel is Z.

## 2.6. SV Calculation

Studies of core dynamics and geomagnetic field behavior require knowledge of SV (time changes of the magnetic field) rather than measurements of the magnetic field itself. Annual means of field data are commonly used, with SV calculated as first differences such that the difference between a given annual mean and the previous annual mean gives the SV at 6 months between the two samples

$$\frac{\partial \mathbf{B}}{\partial t}(t_{m-0.5}) = \mathbf{B}(t_m) - \mathbf{B}(t_{m-1}), \quad (4)$$

with sampling rate  $\Delta t_m = 1$  year and where  $\mathbf{B}$  is the magnetic field and  $t$  is time. However, monthly means are better suited to studies of rapid geomagnetic phenomena due to their higher temporal resolution. Two different types of SV are typically calculated from monthly field averages: annual differences of monthly means (ADMM) and first differences of monthly means (FDMM). ADMM are calculated as the difference between one monthly sample and the same month of the previous year to give the SV 6 months between the two

$$\frac{\partial \mathbf{B}}{\partial t}(t_{n-6}) = \mathbf{B}(t_n) - \mathbf{B}(t_{n-12}), \quad (5)$$

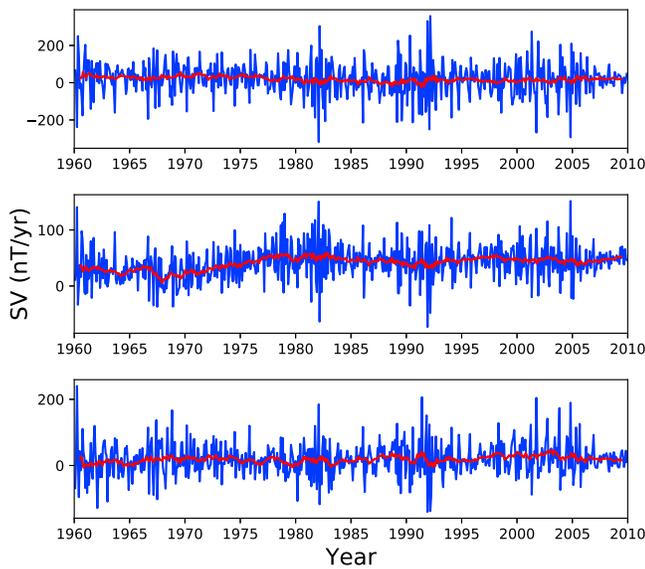
with sampling rate  $\Delta t_n = 1$  month. For example, differencing means at 15 January 2001 and 15 January 2000 gives the SV at 15 July 2000. FDMM are calculated as the difference between two successive monthly field averages

$$\frac{\partial \mathbf{B}}{\partial t}(t_{n-0.5}) = 12[\mathbf{B}(t_n) - \mathbf{B}(t_{n-1})], \quad (6)$$

also with sampling rate  $\Delta t_n = 1$  month, and the factor of 12 is included to convert from nanotesla per month to nanotesla per year. For example, differencing field averages at 15 February 2001 and 15 January 2001 gives the SV at 1 February 2001. ADMM are more frequently used due to their smoothness compared to FDMM, which arises because their calculation is effectively a 12-month running average of the monthly means, both reducing noise and eliminating annual variations. FDMM have higher temporal resolution, but they are much more sensitive to external contamination and other noise sources, see Figure 4 for a comparison of FDMM and ADMM for CLF. MagPySV can also calculate SV using monthly averages and other fractions of years if required, with the user specifying the number of months between successive values (e.g., 12 is used for ADMM and 1 for FDMM).

## 2.7. External Signal Removal Using PCA

A key step in creating usable SV time series is removing external magnetic signals from the observations to leave only time variations of the dynamo-generated magnetic field. Various methods have been used to reduce external signal where possible, such as using only annual means or annual differences of monthly means (e.g., Gillet et al., 2015), data rejection using magnetic indices such as  $K_p$  and  $ap$  (e.g., Hamilton et al., 2015), using nighttime-only data (e.g., Lesur et al., 2015), and removing parameterized models of the external signals from the observations (e.g., Finlay et al., 2016). Here we extend an external noise removal method developed by Wardinski and Holme (2011), which uses PCA. The method uses the differences (“residuals”) between observed SV and that predicted by an observationally constrained internal magnetic field model. The key premise is that residuals provide information about external signals (noise) that are present in the data but not the internal model and that PCA of the residuals covariance matrix gives a proxy for the external noise that is removed from the data. Coherent signals in X, Y, and Z residuals at a single observatory are described by a  $3 \times 3$  covariance matrix (assumed constant through time), whose eigenvectors are used to rotate the residuals into directions of most, intermediate, and least noise. The largest eigenvalue corresponds to the eigendirection with the largest contribution from external noise (the “noisy” direction), and the smallest eigenvalue corresponds to the eigendirection with the lowest contribution from external noise (the “clean” direction). The noisy component residuals for ADMM at the 50 observatories investigated by Wardinski and Holme (2011) are in the north-south plane and have strong zero-lag correlation with the time derivative (annual first differences) of the Dst index, which characterizes magnetic activity of the symmetric equatorial



**Figure 4.** SV at Chambon-la-Forêt, France, calculated as first differences of monthly means (blue) and annual differences of monthly means (red). The top panel is the  $X$  component, middle panel is  $Y$ , and bottom panel is  $Z$ . SV = secular variation.

ring current in the magnetosphere (Sugiura, 1964). The authors used a proxy for external noise in the following noise removal algorithm

$$\hat{R}(t_i) = R(t_i) - \frac{\sum_j^N R(t_j)P(t_j)}{\sum_j^N P(t_j)^2} P(t_i), \quad (7)$$

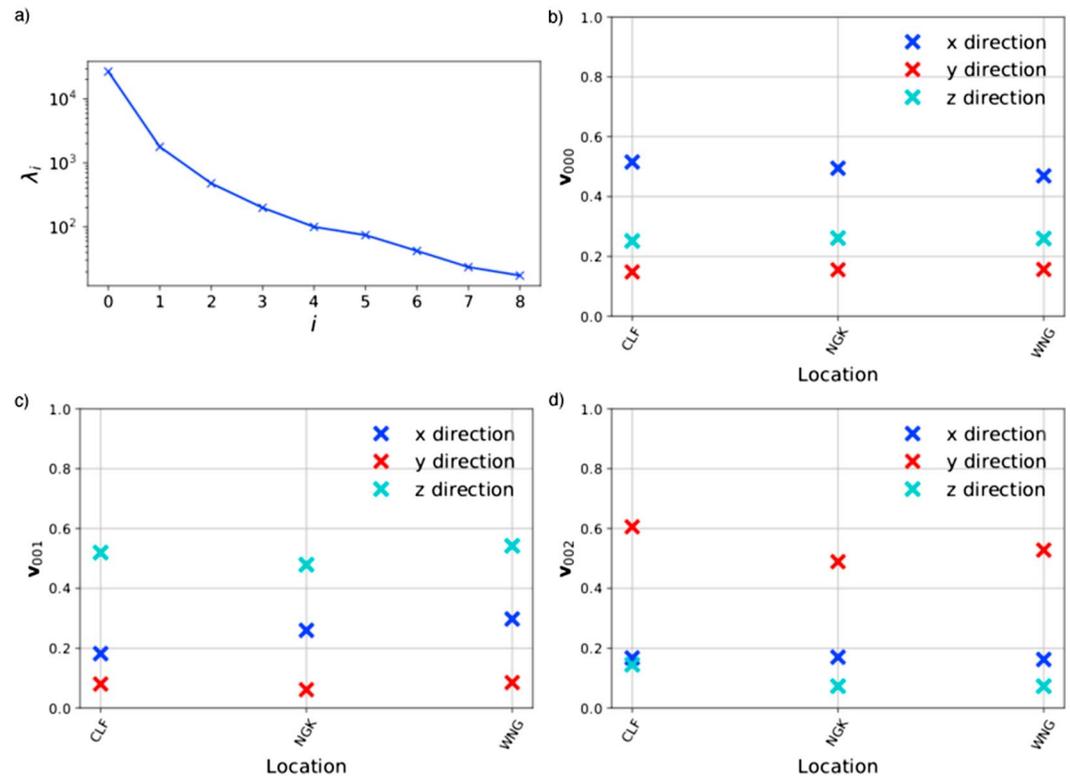
where  $R$  is the noisy component residual,  $\hat{R}$  is the corrected noisy component residual,  $P$  is the proxy for external signal, and the subscripts  $i$  and  $j$  run over all  $N$  time samples. This correction is applied only to residuals in the noisy eigendirection before rotating all residuals back into geographic  $X$ ,  $Y$ , and  $Z$  directions and reforming the SV data (both modeled and unmodeled parts). This yields SV series with reduced external field contamination, with signal having been removed from each geographic component in varying amounts, based on how much it contributed to the noisy direction. In Europe, the noisy direction is predominantly in the  $X$  and  $Z$  directions, due to the geometry of the ring current, which results in the most significant improvements in these directions after denoising, with little change in the  $Y$  component. Wardinski and Holme (2011) found that a stronger correlation is found between the noisy component residuals at different observatories than to the Dst index, leading them to conclude that the noisy component residuals are better able to account for external field variations than the Dst index. This is not surprising as the Dst index is developed using a network of only four observatories and is designed to

monitor the axisymmetric signature of magnetosphere currents (including predominantly the ring current, the tail currents, and the magnetopause Chapman-Ferraro current) rather than all external sources. Therefore, they (and Brown et al., 2013) used the noisy component residual at Niemegek (NGK) observatory, Germany, as the proxy signal in equation (7). NGK was chosen due to having a long unbroken time series that is well documented, and no other observatory could improve on the overall results in that study (NGK itself was corrected by the noisy residual at Chambon-la-Forêt observatory, France).

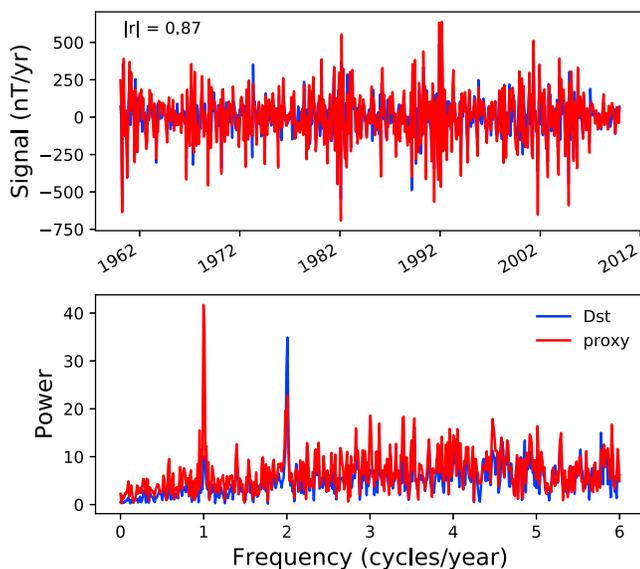
The current work extends this method by considering several observatories simultaneously rather than denoising each individually, though MagPySV performs the original method if only one observatory is specified. This was used to benchmark the current code. Residuals for all observatories of interest (specified by a list of three-digit International Association of Geomagnetism and Aeronomy observatory codes, IAGA) are collated into a single dataframe, whose covariance matrix is of size  $3n \times 3n$  for  $n$  observatories. The eigendirection corresponding to the largest eigenvalue is used as a proxy for the unmodeled external field variations (i.e., noise), which contains contributions from the  $X$ ,  $Y$ , and  $Z$  directions at all  $n$  observatories considered. Depending on the eigenvalue spectrum, the user may choose to use several noisy directions as the proxy signal. If  $m$  eigendirections are used, corresponding to the  $m$  largest eigenvalues, the projected residuals in all  $m$  directions are summed and used as  $P$  in equation (7). Considering several observatories at once permits easy characterization of external signal at groups of nearby observatories, allowing the user to see patterns in the noise. MagPySV outputs figures of the eigenvalue spectrum, the contribution of the geographic components at each observatory to the overall noisy direction, and comparisons of the external signal proxy to the time changes and discrete Fourier transforms of the Dst,  $ap$ , and auroral electrojet (AE) indices, the latter of which characterizes the magnetic signature of the eastward and westward auroral electrojets in the Northern Hemisphere (Davis & Sugiura, 1966). Two applications of this method are presented in the next section, one for Europe and another for northern high latitude observatories.

### 2.7.1. Internal Magnetic Field Models

Several geomagnetic field models are available, each constrained using different data sets that may have been processed in different ways and spanning different time periods. Popular model series include CHAOS (e.g., Finlay et al., 2016; Olsen et al., 2006), gufm (Jackson et al., 2000), COV-OBS (Gillet et al., 2013, 2015), GRIMM (Lesur et al., 2008, 2010, 2015), C<sup>3</sup>FM (Wardinski & Holme, 2006), and Comprehensive Model (Sabaka et al., 2002, 2015). Many core studies focus from 1957 (the International Geophysical Year) onward due to the opening of many geomagnetic observatories and the installation of new instrumentation that provides higher quality measurements. MagPySV uses the COV-OBS model of Gillet et al. (2015), which is



**Figure 5.** (a) Eigenvalue spectrum of the residuals covariance matrix for European observatories and geographic contributions of each European observatory to the eigendirection corresponding to (b) the largest eigenvalue of the residuals covariance matrix, ( $\lambda_0$ , the “noisiest” direction), (c) the second largest eigenvalue of the residuals covariance matrix, ( $\lambda_1$ , the “second noisiest” direction), and (d) the eigendirection corresponding to the third largest eigenvalue of the residuals covariance matrix, ( $\lambda_2$ , the “third noisiest” direction). CLF = Chambon-la-Forêt, France; NGK = Niemegk, Germany; WNG = Wingst, Germany.

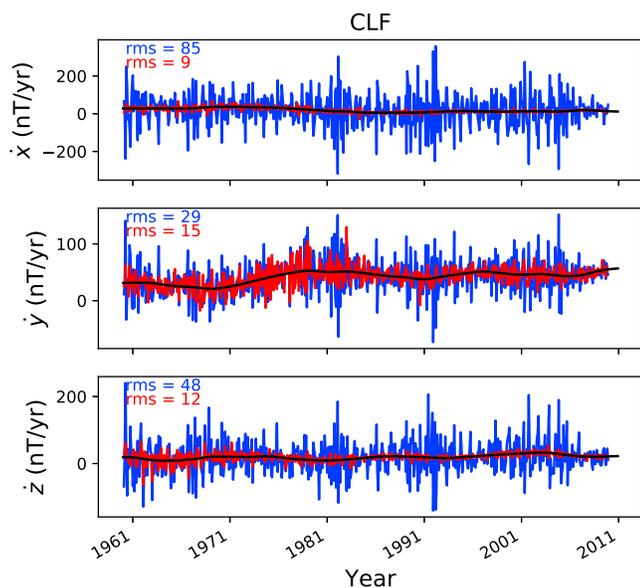


**Figure 6.** Comparison of the proxy used for external signal removal at European observatories and first differences of monthly means of the Dst index. The top panel shows the two time series, and the bottom panel shows their discrete Fourier transforms. Their correlation coefficient is given on the top left.

constrained by observatory and satellite data and spans the period 1840–2020. Our software requires the user to have a compiled executable for the field model; the Fortran source code is available online at <http://www.spacecenter.dk/files/magnetic-models/COV-OBSx1/>, and no modifications to it are necessary in order to run it using our code. MagPySV can obtain the names and locations of all geomagnetic observatories with data stored at WDCE and produce predicted magnetic field and SV time series at those locations for a given date range and frequency, which are then used for denoising and/or plotting. The COV-OBS source code can be easily modified to support spline files for other magnetic field models, in which case, the output will be in the same format and can be read using MagPySV functions. Otherwise, field model output from different codes can be read and manipulated using standard Python commands.

### 2.7.2. Treatment of Missing Data

The built in PCA method of Python’s Scikit-learn package (Pedregosa et al., 2011) uses singular value decomposition, which is more efficient than calculating the covariance matrix of the residuals and then finding its eigenvalues/vectors. However, this method cannot be used when any data are missing, and the user must either infill data or remove rows with NaN values. Although MagPySV does include a denoising function that infills missing data using imputation, this is usually not appropriate because many series have large gaps (several months or years), or the period of interest may begin several years prior to the opening of some



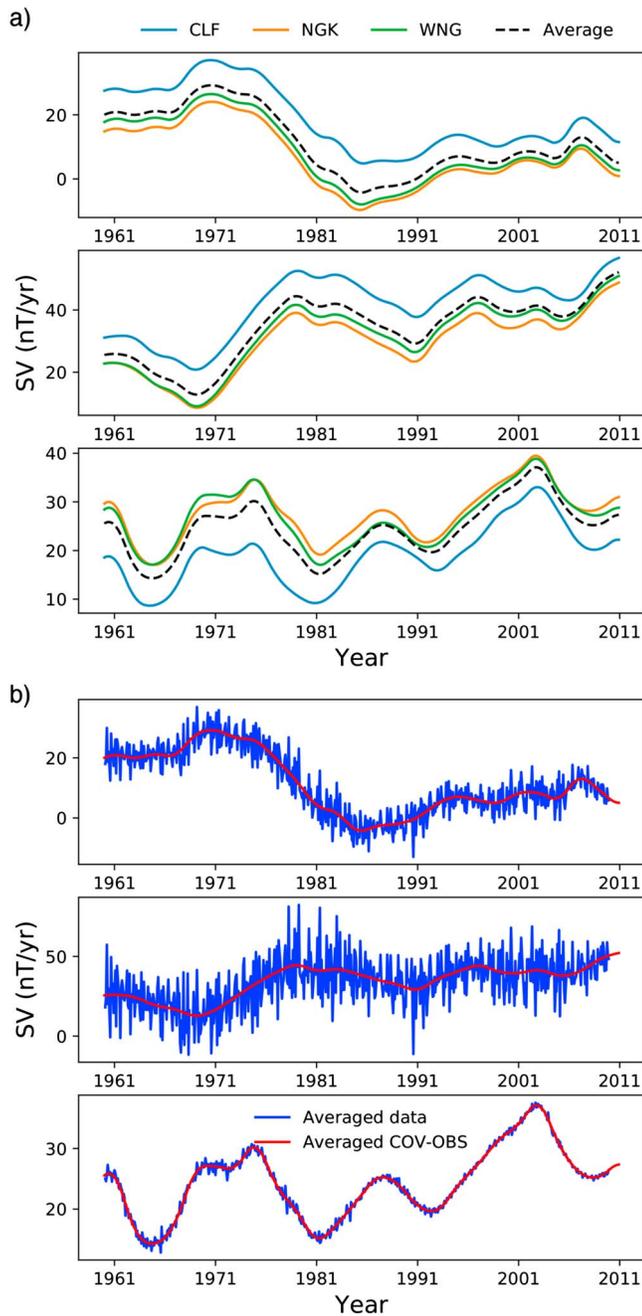
**Figure 7.** First differences of monthly means at Chambon-la-Forêt before (blue) and after (red) denoising with residuals in the two noisiest eigendirections used as the proxy for unmodeled external signal.

observatories among those considered. Dropping rows containing one or more NaN values means removing all data for that date (a month if using FDMM or ADMM or year if using annual means) even if the value is missing at only one observatory out of all those considered. Since there are many gaps in different observatories at different times, this potentially results in discarding the majority of usable data. The recommended implemented method uses masked arrays to calculate covariance matrix of non-NaN values, allowing the use of “gappy” series without discarding useful information and data from observatories whose series do not cover the entire period of interest.

### 3. Case Study I: European Observatories

Europe has a dense network of high quality magnetic observatories, many with several decades of continuous data. Consequently, SV has been well-studied in this region using ADMM, and several geomagnetic jerks have been first identified there. However, no studies using FDMM have been reported due to their high noise content. In this case study, we use MagPySV to denoise FDMM between 1960 and 2010 for a group of three European observatories simultaneously (CLF, France; NGK, Germany; and Wingst, Germany) to produce higher-resolution SV series than previously investigated. In particular, we are interested in whether the internal field model is able to reproduce sharp changes in the data at jerk occurrence

times and whether the new series show evidence of the alternative jerk morphology proposed by Holme and de Viron (2013). On the basis of a discontinuity in length-of-day series, they proposed a discontinuity in the SV itself for the 2003.5 event rather than the well-established definition of a jerk being a change in SV trend (the secular acceleration). Figure 5a shows the eigenvalue spectrum of the residuals covariance matrix when all observatories are considered together, having performed all initial processing steps except applying an  $\alpha p$  threshold and removing outliers. Note that Python assigns indices from 0 rather than 1, so that, for example, the eigenanalysis of the covariance matrix for three observatories yields nine eigenvalues (and corresponding eigenvectors) that are numbered from 0 to 8. Since this paper accompanies Python software, the figures presented here follow Python numbering convention. The largest eigenvalue,  $\lambda_0$ , corresponds to the noisiest eigendirection,  $\mathbf{v}_0$ ; that is, the direction contributing the most to the unmodeled external signals. This eigendirection is predominantly in the X direction, with some Z, and the second noisiest direction,  $\mathbf{v}_1$ , corresponding to  $\lambda_1$ , is predominantly in the Z, with some X (see Figures 5b and 5c). As previously reported in Wardinski and Holme (2011) for single observatory denoising, these geographic components correspond to the direction of the magnetic signature of the magnetospheric ring current relative to Europe and are likely attributable to that external source. In the third noisiest direction, corresponding to  $\lambda_2$ , a dominant Y contribution occurs at all locations (Figure 5d). A coherent signal in the residuals could arise via two sources; first, an external signal that is present in data but not the internal field model (noise), or second, an internal field variation that is not captured by the internal field model, possibly because such smoothed models sometimes cannot replicate the sharpest temporal changes in the data, such as jerks. In separate analyses (not shown here), we investigated several further groups of observatories in Europe, North America, and Asia, considering both regional and global observatory combinations, and found that this dominant Y contribution to  $\mathbf{v}_2$  is present across all of Europe but not in the other geographic regions. Since many jerks are most clearly seen in the Y component, this signal may be of internal origin, appearing coherently across European observatories due to their small geographic spread. However, a Fourier analysis of the time series in that direction (the discrete Fourier transform) does show a significant semiannual component, from which we infer that this direction is at least partly of external origin perhaps related to non-axisymmetric magnetospheric ring currents. Since we are not able to separate the internal and external parts of  $\mathbf{v}_2$ , we leave it in the data and use only the residuals in eigendirections  $\mathbf{v}_0$  and  $\mathbf{v}_1$  to denoise in this example. Note that the dominant Z signal in  $\mathbf{v}_1$  may also be of partly internal origin; see Figure 6 for a comparison of our external signal proxy with Dst, which shows a strong correlation coefficient and that both signals contain strong annual and semiannual components. We use time changes (FDMM to match the used SV data) of the definitive Dst index from WDC Kyoto (available at <http://wdc.kugi.kyoto-u.ac.jp/dstae/index.html>) for



**Figure 8.** (a) Field model predictions for the  $X$  (top),  $Y$  (middle), and  $Z$  (bottom) SV components at CLF, NGK, and WNG, along with the averaged field model over these three observatories (black-dashed line) and (b) denoised  $X$  (top),  $Y$  (middle), and  $Z$  (bottom) SV components averaged over CLF, NGK, and WNG, along with the averaged field model (red). CLF = Chambon-la-Forêt, France; NGK = Niemegk, Germany; WNG = Wingst, Germany; SV = secular variation.

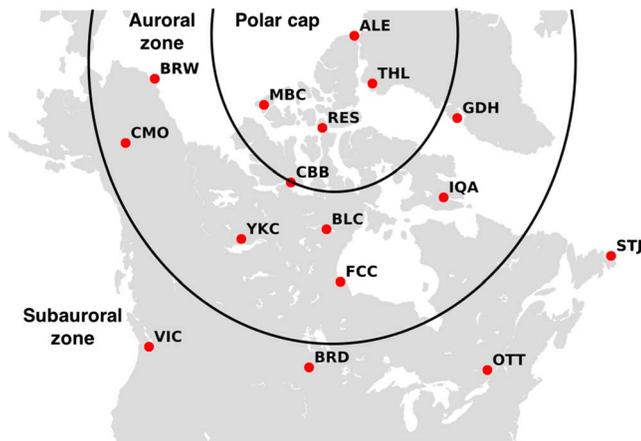
comparison with our external noise proxies; we have also tried using the Dcx index (corrected, extended Dst available at <http://dcx.oulu.fi/>; Mursula et al., 2008) but found much weaker correlations than with the original Dst index. Presumably, this is because significant annual and semiannual variations are present in our noise proxies; these variations have been removed from the Dcx index but not Dst. The denoising (using the projected residuals in the two noisiest eigendirections as the proxy for external signal in (7) results in SV series with much reduced external contamination; see Figure 7, for example, which shows the SV for CLF observatory pre-denoising (blue) and postdenoising (red) and the root-mean-square of the residuals for comparison. The  $Z$ , and to a lesser extent  $X$ , components show much improvement (in the sense that high magnitude, rapid variation due to external sources has been removed while the internal field variation remains), as they contributed most to the noisy eigendirections. See also Shore et al. (2016), who used PCA of observatory hourly means during geomagnetically quiet days (at middle to low latitudes) to extract the dominant spatiotemporal patterns of long-period external fields.

Given their close geographic proximity, it is of interest to consider whether averaging data over these three locations can further reduce noise. The model predictions for SV are similar in the three locations (Figure 8a) and so an average is suitably representative of the individual locations. Averaging the data yields the SV series shown in Figure 8b, which clearly show that the field model very closely follows the general trend of the data, particularly in the  $Z$  and  $X$  components, which is unsurprising as the dense European data are used to constrain the field model in the first place. In  $Y$ , the data permit sharper temporal changes than present in the model, particularly at previously determined jerk occurrence times (Brown et al., 2013), which perhaps explains the  $Y$  dominance in the  $\mathbf{v}_2$  direction that was left in the data. However, even with higher-resolution series than previously studied, we see no evidence of a discontinuity in SV itself at jerk times (as proposed by Holme & de Viron, 2013) rather than in its time derivative, the secular acceleration, for the 2003.5 event.

#### 4. Case Study II: High Latitude Observatories

High latitude observatories provide vital information about external field variations, in particular the field-aligned currents, aurorae, and dynamics of the magnetosphere, but are often neglected from internal magnetic field studies due to their high external signal content. However, observationally constrained field models have shown marked differences in SV at high and low latitudes, such as four persistent quasi-stationary high latitude equatorially symmetric flux lobes (Bloxham & Gubbins, 1985), which are believed to be a surface manifestation of rotationally aligned convection rolls in the outer core. Livermore et al. (2017) used recent satellite data (Finlay et al., 2016; Olsen et al., 2014) to infer the presence of an accelerating jet on the tangent cylinder due to the observed acceleration of the Canadian and Siberian flux lobes. The tangent cylinder is the imaginary cylinder in the outer core that just surrounds the inner core and is parallel to the axis of rotation and manifests in the observable geomagnetic field

at high latitudes. There is evidence that the tangent cylinder is the source region for torsional waves in the outer core, observations of which allow the estimation of magnetic field strength inside the core (Gillet et al., 2010). Thus, the high latitudes are of great interest to core dynamicists, but efforts to tie observations to their underlying physical mechanisms are hindered by the lack of long SV series in these regions because satellite missions provide excellent data quality and resolution but are temporally discontinuous and of limited



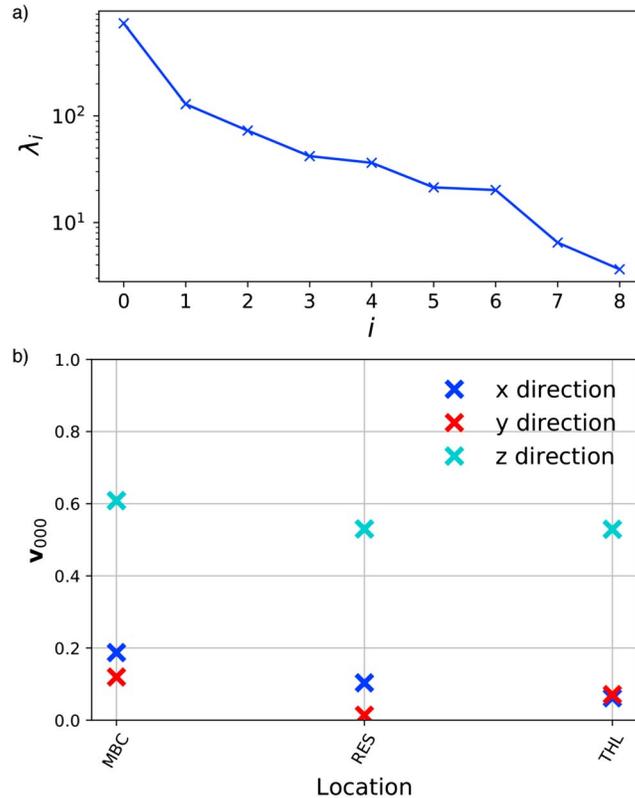
**Figure 9.** Map of selected magnetic observatories located within the polar cap, auroral zone, and subauroral zone. The black lines denote approximate boundaries between the three high latitude regions after Nikitina et al. (2016) Figure 1. THL = Thule, Greenland; MBC = Mould Bay, Canada; RES = Resolute Bay, Canada; CBB = Cambridge Bay, Canada; BRW = Barrow, Alaska; YKC = Yellowknife, Canada; BLC = Baker Lake, Canada; GDH = Qeqertarsuaq (Godhavn), Greenland; VIC = Victoria, Canada; OTT = Ottawa, Canada; STJ = St. John's, Canada; ALE = Alert, Canada; CMO = College, USA; FCC = Fort Churchill, Canada; IQA = Iqaluit, Canada; BRD = Brandon, Canada.

duration, and ground-based observatories have longer series but are subject to often overwhelming contamination. In this case study, we investigate SV at three clusters of observatories, grouped in different zones according to their geomagnetic latitude. The zones are referred to as the polar cap, the auroral zone, and the subauroral zone (Figure 9).

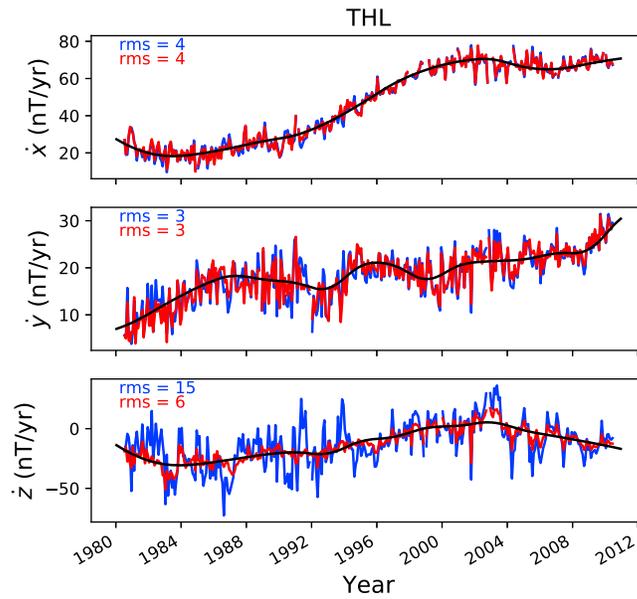
#### 4.1. Polar Cap

We consider three observatories in the polar cap: Mould Bay, Canada; Resolute Bay, Canada; and Thule, Greenland (THL); see Figure 9 for their locations. The eigenvalue spectrum shows a change of slope after the largest eigenvalue  $\lambda_0$  (Figure 10a), and the corresponding noisy eigendirection is predominantly Z, indicating that noise is vertical near geomagnetic pole (Figure 10b). Denoising using residuals in direction  $\mathbf{v}_0$  results in significant improvements in the Z SV component, see Figure 11 for an example at THL, but the other components remain the same as before. In order to see improvements in all components, we need to use residuals in additional eigendirections in the noise proxy. However, we did not see any coherent patterns in the next few eigendirections, so a more detailed analysis of residuals is needed here. Being mostly vertical, we found no strong correlations between the noise proxy and time variations (ADMM to match the used SV data) of any of the geomagnetic indices we consider (Dst [correlation coefficient 0.55],  $ap$  [0.54], and AE [0.60]), which is unsurprising as none of these indices describes activity over the polar caps. Although the

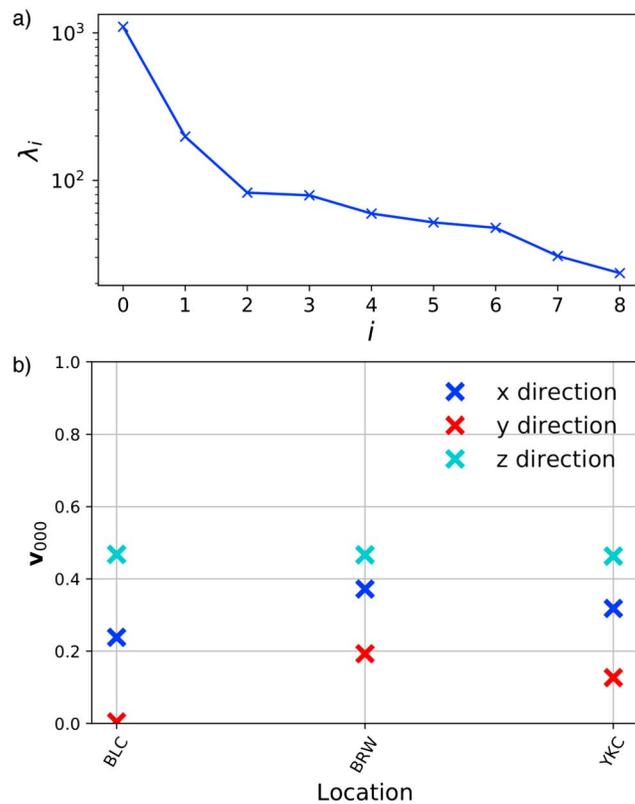
Polar Cap Index (Troshichev et al., 1988) is designed for this purpose, we did not consider it in the present analysis because we include THL observatory, whose data are used to derive the Polar Cap North Index. In a



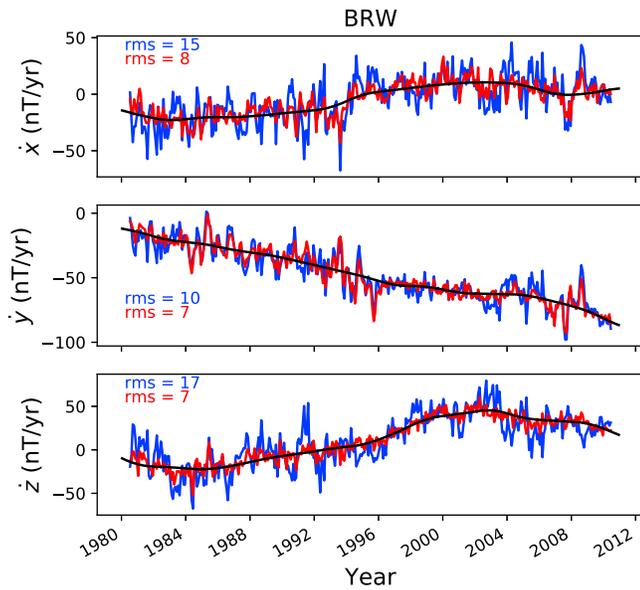
**Figure 10.** (a) Eigenvalue spectrum of the residuals covariance matrix for polar observatories and (b) geographic contributions of each polar observatory to the eigendirection corresponding to the largest eigenvalue of the residuals covariance matrix, ( $\lambda_0$ , the “noisiest” direction). MBC = Mould Bay, Canada; RES = Resolute Bay, Canada; THL = Thule, Greenland.



**Figure 11.** Annual differences of monthly means secular variation at THL before (blue) and after (red) denoising with residuals in the noisiest eigendirection used as the proxy for unmodeled external signal. THL = Thule, Greenland.



**Figure 12.** (a) Eigenvalue spectrum of the residuals covariance matrix for auroral observatories and (b) geographic contributions of each auroral observatory to the eigendirection corresponding to the largest eigenvalue of the residuals covariance matrix, ( $\lambda_0$ , the “noisiest” direction). BLC = Baker Lake, Canada; BRW = Barrow, Alaska; YKC = Yellowknife, Canada.

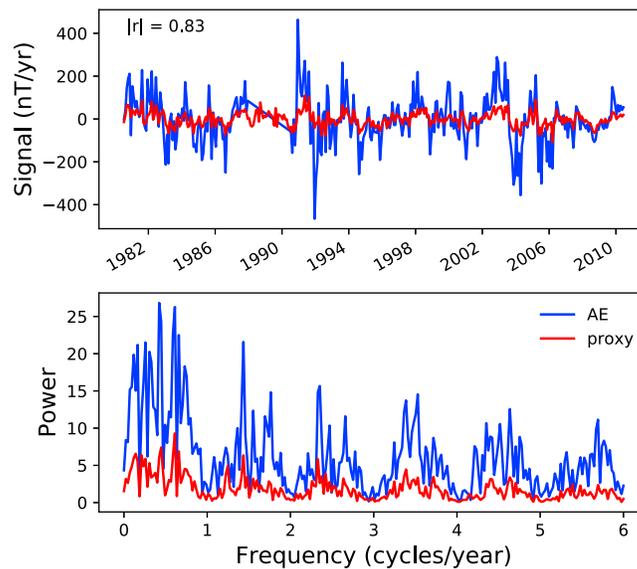


**Figure 13.** Annual differences of monthly means at BRW before (blue) and after (red) denoising with residuals in the noisiest eigendirection used as the proxy for unmodeled external signal. BRW = Barrow, Alaska.

Barrow, Alaska, in Figure 13 shows improvements in all three components, though Y is the least improved. The noisy direction has a relatively strong correlation with the AE index (coefficient is 0.83; Figure 14) and a slightly weaker correlation with Dst (0.80). We have tried including Cambridge Bay, Canada, observatory in the auroral observatories, which lies on the boundary between the polar cap and auroral zone according to Figure 9, but found that its noise is predominantly vertical, which is more consistent with the polar cap than this region. Again, this indicates that PCA is a useful method for determining to which high latitude region an observatory belongs based on its external signal content rather than only its location.

#### 4.3. Subauroral Zone

We consider three observatories in the subauroral zone: Ottawa, Canada; St. John's, Canada; and Victoria, Canada; see Figure 9 for their locations. The eigenvalue spectrum has a change in slope at  $\lambda_1$  (Figure 15a),

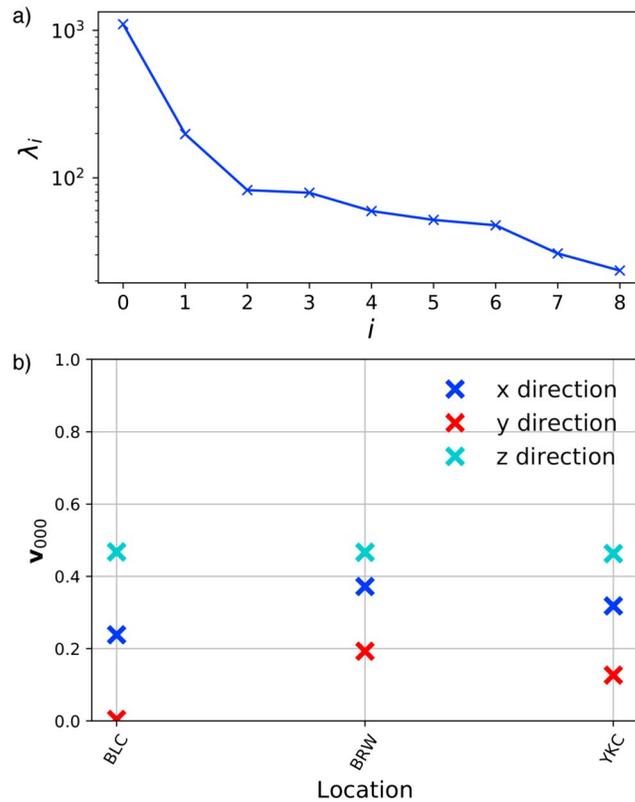


**Figure 14.** Comparison of the proxy used for external signal removal at auroral observatories and annual differences of monthly means of the AE index. The top panel shows the two time series, and the bottom panel shows their discrete Fourier transforms. Their correlation coefficient is given on the top left.

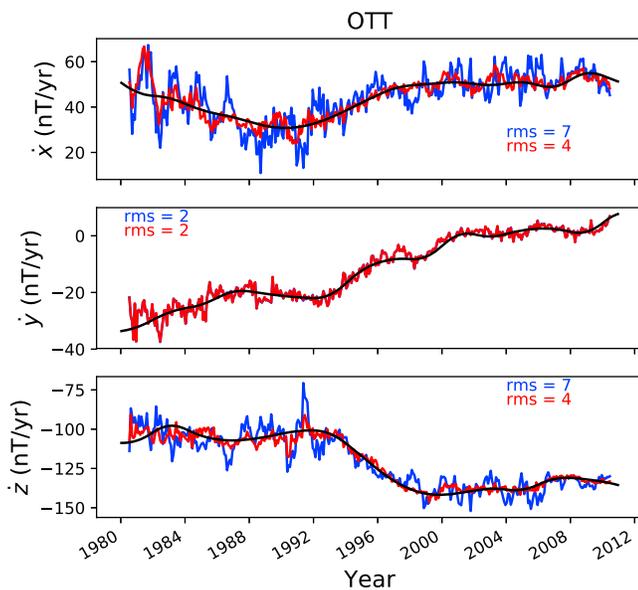
separate analysis (not shown), we also considered Qeqertarsuaq (Godhavn), Greenland, observatory (GDH) as part of the polar cap cluster of observatories, but we found that its noise content is different from those presented here. However, it is similar to that of the auroral observatories described in the next section. Since its location is close to the approximate boundary between the polar cap and auroral zone in Figure 9, and the real boundary is difficult to determine precisely, GDH observatory could be in either of these geomagnetic zones. As our PCA reveals very different external signal directions in these two regions, and those at GDH are consistent with the auroral observatories, we argue that GDH belongs to that high latitude region. We further suggest that this method could be used to better categorize other observatories that lie near the boundaries between different geomagnetic regions.

#### 4.2. Auroral Zone

We consider three observatories in the auroral zone: Baker Lake, Canada; Barrow, Alaska; and Yellowknife, Canada; see Figure 9 for their locations. The eigenvalue spectrum has a clear change in slope at  $\lambda_2$ , separating a steep slope for the first two eigenvalues and a shallow slope for the rest (Figure 12a). The signals in the two noisiest eigendirections are mostly horizontal (X and Z), see Figure 12b for an example. Using only residuals in  $\mathbf{v}_0$  and using both  $\mathbf{v}_0$  and  $\mathbf{v}_1$  as the external signal proxy for denoising gives very similar results for this group of observatories. We therefore show results using the noisiest direction only, for example, the SV at Barrow, Alaska, in Figure 13 shows improvements in all three components, though Y is the least improved.



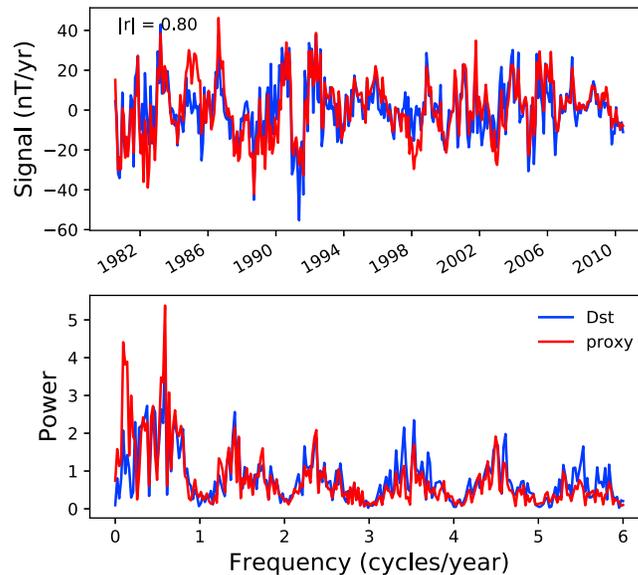
**Figure 15.** (a) Eigenvalue spectrum of the residuals covariance matrix for subauroral observatories and (b) geographic contributions of each subauroral observatory to the eigendirection corresponding to the largest eigenvalue of the residuals covariance matrix,  $\lambda_0$ , the “noisiest” direction). OTT = Ottawa, Canada; STJ = St. John’s, Canada; VIC = Victoria, Canada.



**Figure 16.** Annual differences of monthly means at OTT before (blue) and after (red) denoising with residuals in the noisiest eigendirection used as the proxy for unmodeled external signal. OTT = Ottawa, Canada.

and the residuals in the noisy eigendirection are almost entirely horizontal and equally split between  $X$  and  $Z$  (Figure 15b). We use the residuals in  $\mathbf{v}_0$  as the external signal proxy, see Figure 16 for an example of the results at Ottawa, Canada. The noise proxy has a relatively strong correlation with the Dst index (coefficient 0.80; Figure 17), though this is weaker than the correlation between the European noisy direction and Dst despite the similar latitudes.

In these two case studies, we have shown that ADMM SV series can be significantly improved upon, and usable high-resolution FDMM series constructed, by implementing the proposed denoising method, though not all components are improved to the same degree. We have presented cleaned high-resolution SV series for European observatories and cleaned data for high northern latitude observatories. Cleaner results could be obtained by removing more eigendirections when denoising, though care must be taken to determine that residuals in these directions are of external origin; otherwise, the internal field variations of interest will be also removed. However, the simple analyses presented in these case studies illustrate the key point; we are able to characterize noise in different geographic regions using this method, and MagPySV provides an easy and reproducible means to do this. We are able to easily identify observatories at which the noise behaves similarly, such as in different geomagnetic latitude bands, and have found that the best denoising results are obtained when considering groups of observatories with similar noise profiles.



**Figure 17.** Comparison of the proxy used for external signal removal at subauroral observatories and annual differences of monthly means of the Dst index. The top panel shows the two time series, and the bottom panel shows their discrete Fourier transforms. Their correlation coefficient is given on the top left.

## 5. Code Implementation and QC

The MagPySV code is under version control through GitHub and is freely available at <https://github.com/gracecox/MagPySV> under the MIT/X11 license. The latest working version is in the master branch, and new features are implemented in various development branches. For future reference, the specific MagPySV version that accompanies this paper may be accessed using its digital object identifier at <https://doi.org/10.5281/zenodo.1340640>. Parts of the code have associated unit and functional tests, which are run via continuous integration with Travis upon each push to the master branch. The documentation is hosted online at <http://magpysv.readthedocs.io/en/latest/>; the current status of latest changes is displayed on documentation homepage (whether build is passing or failing, etc). The code is available on PyPI and may be installed using pip (recommended) or alternatively from source via git clone. Two Jupyter notebook tutorials of MagPySV's workflow are found in a separate GitHub repository at <https://github.com/gracecox/MagPySV-examples>, which reproduces all figures shown in this paper except the flowchart in Figure 1. Detailed installation instructions are on the GitHub project homepage and on the documentation homepage.

The code is written in Python language version 3.x, which is mostly, but not completely, backward compatible with Python 2.x languages and a set of third party libraries including NumPy (van der Walt et al., 2011), SciPy (Jones et al., 2001), Pandas (McKinney, 2010), Scikit-learn (Pedregosa et al., 2011), Matplotlib (Hunter, 2007), and Jupyter (Pérez & Granger, 2007). The coding style is mostly procedural programming rather than object-oriented to make implementation of desired functionality more transparent and easier to follow. It complies with PEP8 recommendations for Python programming style, which are designed to improve readability, consistency of code between different developers, and ease of maintenance. Adherence to Python programming standards is checked using Codacy, an automated code analysis/quality tool that gives code quality metrics on each push to the repository. The package is itself a library (a collection of modules) so does not include a graphical user interface. The code may be run on the command line or via interactive Jupyter notebooks, the latter of which contain interspersed text, code snippets, and code output such as figures. Each module/function is documented using sphinx-docs formatted doc strings that specify the purpose of the code and any inputs/outputs, which can be accessed through the documentation or the help command in Python. We would welcome feedback and/or suggestions for future extensions to the software. Any member of the community can contribute to this project by forking (downloading) the MagPySV repository from GitHub, making their changes to the code and then requesting that they be merged with the master branch via a pull request.

The library of code to access the BGS GDP is also written in Python 3.x and follows the same ethos of coding style and unit testing as MagPySV. The code can function as a standalone package, available from PyPI through pip as `gmdata_webinterface`, and is installed automatically as a dependency of MagPySV. The source is also open under an MIT/X11 license and version controlled, available on GitHub at [https://github.com/willjbrown88/geomag\\_wdc\\_web\\_app\\_interface](https://github.com/willjbrown88/geomag_wdc_web_app_interface). We again welcome comments and contributions as we look to expand the functionality in future updates.

## 6. Summary

We have presented MagPySV, a Python package for obtaining magnetic observatory hourly means from WDC Edinburgh, processing the raw data, and removing external magnetic field contamination. Much of this data processing is currently performed using closed-source piecemeal codes or on an ad hoc basis, which impedes efforts to reproduce the data sets underlying publications. The software provides a consistent and automated method for creating time series suitable for studying internal field variations and their underlying core processes. The eigensystem method employed permits use of all data (except outliers) rather than discarding potentially useful data to reduce noise (e.g., by using only local nighttime data) or by performing data selection using a geomagnetic index, which may be inappropriate because no single geomagnetic index is designed to capture all external field sources in all geographic regions. Also, such indices are typically created using only a handful of observatories and are therefore biased to particular geographic regions. We have presented two case studies of denoised data: first in Europe, which has a dense observatory network and has been well studied using ADMM. We have produced cleaned high-resolution SV series (FDMM) but found no evidence for geomagnetic jerks manifesting as a discontinuity in SV (as proposed by Holme & de Viron, 2013) rather than in secular acceleration. In the second case study, we have shown that the noise content at high northern latitudes is different in three zones corresponding to geomagnetic latitude (the polar cap, auroral zone, and subauroral zone). Considering groups of observatories whose noise content is similar produces the best results and allows us to make better use of available data for studying core dynamics, especially in or around the tangent cylinder. We hope this software will be of use to the geomagnetic community and welcome their contributions. The code is of potential interest to researchers who wish to make their work reproducible, those investigating rapid geomagnetic SV (e.g., jerks) who would like data at higher temporal resolution, and those who invert observatory data to investigate core flow. It may also appeal to core dynamicists and numerical modelers who do not usually work with geomagnetic observatory data due to the time and expertise involved but who would like to draw comparisons between their models and observations given easy access to the data. Since parameterized models of external magnetic field variations do not capture all signals due to our incomplete knowledge of the physical processes involved, PCA offers an advantage in that it is able to better characterize and extract noise from the data. Of course, though we have referred to external magnetic field variations as noise from the point of view of internal magnetic field studies, they are the subject of interest in other disciplines. Our external signal proxies obtained through PCA provide information about external signals that is not gained through other methods, which has obvious applications to external field studies. We are able to look in detail at residuals projected into the different eigendirections, which helps to identify which external sources are contributing the most at different geomagnetic latitudes. Additionally, we envisage immediate applications of MagPySV to electromagnetic induction studies because the noise proxies can be used as the external source rather than the simplified (often axisymmetric) spherical harmonic parameterized models currently employed.

## References

- Barnett, V., & Lewis, T. (1974). *Outliers in statistical data*: Wiley.
- Bloxham, J., & Gubbins, D. (1985). The secular variation of Earth's magnetic field. *Nature*, 317, 777–781.
- Brown, W., Mound, J., & Livermore, P. (2013). Jerks abound: An analysis of geomagnetic observatory data from 1957 to 2008. *Physics of the Earth and Planetary Interiors*, 223, 62–76. <https://doi.org/10.1016/j.pepi.2013.06.001>
- Bureau Central De Magnétisme Terrestre (1921). French global network of geomagnetic observatories (version 1). <https://doi.org/10.18715/bcmt.mag.def>
- Chulliat, A., & Telali, K. (2007). World monthly means database project. *Publications of the Institute of Geophysics, Polish Academy of Sciences*, 99, 398.
- Clarke, E., Baillie, O., Reay, S. J., & Turbitt, C. W. (2013). A method for the near real-time production of quasi-definitive magnetic observatory data. *Earth, Planets and Space*, 65(11), 16.
- Davis, T. N., & Sugiura, M. (1966). Auroral electrojet activity index *ae* and its universal time variations. *Journal of Geophysical Research*, 71(3), 785–801.

### Acknowledgments

We thank the national institutes that support geomagnetic observatories, INTERMAGNET for promoting high standards of magnetic observatory practice ([www.intermagnet.org](http://www.intermagnet.org)), and the WDC for Geomagnetism at BGS, Edinburgh. The specific data sets used for the Dst and AE indices may be accessed through World Data Center for Geomagnetism Kyoto et al. (2015a) and World Data Center for Geomagnetism Kyoto et al. (2015b). We have used definitive BCMT data, which are available through Bureau Central De Magnétisme Terrestre (1921). The magnetic observatory data used in this paper are available through the BGS Geomagnetism Data Portal at <http://wdc.bgs.ac.uk/dataportal/>. MagPySV's example Jupyter notebooks can be used to download these data and reproduce the presented results. We thank Susan Macmillan for providing advice on data processing and the file containing all recorded baseline discontinuities. The authors thank two anonymous reviewers for comments and suggestions that helped to improve the manuscript. G. A. C. thanks Andrew Walker for advice on software development and coding in Python. G. A. C. and R. H. were supported by NERC grant NE/M012190/1.

- Feng, Y., Holme, R., Cox, G. A., & Jiang, Y. (2018). The geomagnetic jerk of 2003.5—characterisation with regional observatory secular variation data. *Physics of the Earth and Planetary Interiors*, 278, 47–58.
- Finlay, C. C., Olsen, N., Kotsiaros, S., Gillet, N., & Toffner-Clausen, L. (2016). Recent geomagnetic secular variation from Swarm. *Earth, Planets and Space*, 68(1), 1–18.
- Gillet, N., Barrois, O., & Finlay, C. C. (2015). Stochastic forecasting of the geomagnetic field from the COV-OBS.x1 geomagnetic field model, and candidate models for igrf-12. *Earth, Planets and Space*, 67(1), 71.
- Gillet, N., Jault, D., Canet, E., & Fournier, A. (2010). Fast torsional waves and strong magnetic field within the Earth's core. *Nature*, 465(74), 74–77.
- Gillet, N., Jault, D., Finlay, C. C., & Olsen, N. (2013). Stochastic modeling of the Earth's magnetic field: Inversion for covariances over the observatory era. *Geochemistry, Geophysics, Geosystems*, 14, 766–786. <https://doi.org/10.1002/ggge.20041>
- Hamilton, B., Ridley, V. A., Beggan, C. D., & Macmillan, S. (2015). The BGS magnetic field candidate models for the 12th generation IGRF. *Earth Planets Space*, 67(1), 69. <https://doi.org/10.1186/s40623-015-0227-x>
- Hampel, F. R. (1974). The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346), 383–393.
- Hodge, V., & Austin, J. (2004). A survey of outlier detection methodologies. *Artificial Intelligence Review*, 22(2), 85–126.
- Holme, R., & de Viron, O. (2013). Characterization and implications of intradecadal variations in length of day. *Nature*, 499, 202–204.
- Hulot, G., Sabaka, T., & Olsen, N. (2007). The present field. In *Treatise on geophysics* (Vol. 5). Amsterdam: Elsevier.
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing In Science & Engineering*, 9(3), 90–95.
- Jackson, A., Jonkers, A. R. T., & Walker, M. R. (2000). Four centuries of geomagnetic secular variation from historical records. *Philosophical Transactions of the Royal Society A*, 358, 957–990.
- Jones, E., Oliphant, T., & Peterson, P. (2001). SciPy: Open source scientific tools for Python. accessed: 2017-03-01.
- Lesur, V., Rother, M., Wardinski, I., Schachtschneider, R., Hamoudi, M., & Chambodut, A. (2015). Parent magnetic field models for the IGRF-12 GFZ candidates. *Earth Planets Space*, 67(1), 1–15.
- Lesur, V., Wardinski, I., Hamoudi, M., & Rother, M. (2010). The second generation of the GFZ reference internal magnetic model: GRIMM-2. *Earth, Planets and Space*, 62(10), 6.
- Lesur, V., Wardinski, I., Rother, M., & Manda, M. (2008). GRIMM: The GFZ reference internal magnetic model based on vector satellite and observatory data. *Geophysical Journal International*, 173(2), 382–394.
- Livermore, P. W., Hollerbach, R., & Finlay, C. C. (2017). An accelerating high-latitude jet in Earth's core. *Nature Geoscience*, 10(1), 62.
- Macmillan, S., & Olsen, N. (2013). Observatory data and the Swarm mission. *Earth, Planets and Space (Online)*, 65(11), 1355.
- Manda, M., Holme, R., Pais, A., Pinheiro, K., Jackson, A., & Verbanac, G. (2010). Geomagnetic jerks: Rapid core field variations and core dynamics. *Space Science Reviews*, 155(1–4), 147–175.
- McKinney, W. (2010). Data structures for statistical computing in Python. In *Proceedings of the 9th Python in Science Conference* (Vol. 445, pp. 51–56). SciPy Austin, TX.
- Mursula, K., Holappa, L., & Karinen, A. (2008). Correct normalization of the Dst index. *Astrophysics and Space Sciences Transactions*, 4(2), 41–45.
- Nikitina, L., Trichtchenko, L., & Boteler, D. (2016). Assessment of extreme values in geomagnetic and geoelectric field variations for Canada. *Space Weather*, 14, 481–494. <https://doi.org/10.1002/2016SW001386>
- Olsen, N., Lühr, H., Finlay, C. C., Sabaka, T. J., Michaelis, I., Rauberg, J., & Toffner-Clausen, L. (2014). The CHAOS-4 geomagnetic field model. *Geophysical Journal International*, 197(2), 815–827.
- Olsen, N., Lühr, H., Sabaka, T. J., Manda, M., Rother, M., Toffner-Clausen, L., et al. (2006). CHAOS—A model of the Earth's magnetic field derived from CHAMP, Ørsted, and SAC-C magnetic satellite data. *Geophysical Journal International*, 166, 67–75.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Peltier, A., & Chulliat, A. (2010). On the feasibility of promptly producing quasi-definitive magnetic observatory data. *Earth, Planets and Space*, 62(2), e5–e8.
- Pérez, F., & Granger, B. E. (2007). IPython: A system for interactive scientific computing. *Computing in Science & Engineering*, 9(3), 21–29.
- Rangarajan, G. K. (1989). Geomagnetism. In *Chap. indices of geomagnetic activity* (Vol. 3, pp. 323–384). London: Academic Press.
- Sabaka, T., Olsen, N., & Langel, R. (2002). A comprehensive model of the quiet-time, near-Earth magnetic field: Phase 3. *Geophysical Journal International*, 151(1), 32–68.
- Sabaka, T. J., Olsen, N., Tyler, R. H., & Kuvshinov, A. (2015). Cm5, a pre-Swarm comprehensive geomagnetic field model derived from over 12 yr of CHAMP, Ørsted, SAC-C and observatory data. *Geophysical Journal International*, 200(3), 1596–1626.
- Shore, R., Whaler, K., Macmillan, S., Beggan, C., Velínský, J., & Olsen, N. (2016). Decadal period external magnetic field variations determined via eigenanalysis. *Journal of Geophysical Research: Space Physics*, 121, 5172–5184. <https://doi.org/10.1002/2015JA022066>
- Sugiura, M. (1964). Hourly values of equatorial Dst for the IGY. *Annals of the International Geophysical Year*, 35, 7–45.
- Troshichev, O., Andrezen, V., Vennerstrøm, S., & Friis-Christensen, E. (1988). Magnetic activity in the polar cap—A new index. *Planetary and Space Science*, 36(11), 1095–1102.
- van der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The NumPy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2), 22–30.
- Wardinski, I., & Holme, R. (2006). A time-dependent model of the Earth's magnetic field and its secular variation for the period 1980–2000. *Journal of Geophysical Research*, 111, B12101. <https://doi.org/10.1029/2006JB004401>
- Wardinski, I., & Holme, R. (2011). Signal from noise in geomagnetic field modelling: Denoising data for secular variation studies. *Geophysical Journal International*, 185(2), 653–662.
- World Data Center for Geomagnetism Kyoto, Nose, M., Iyemori, T., Sugiura, M., & Kamei, T. (2015a). Geomagnetic Dst index (version 1). <https://doi.org/10.17593/14515-74000>
- World Data Center for Geomagnetism Kyoto, Nose, M., Iyemori, T., Sugiura, M., & Kamei, T. (2015b). Geomagnetic AE index (version 1). <https://doi.org/10.17593/15031-54800>