

Geomagnetism Program

Adjusted Geomagnetic Data—Theoretical Basis and Validation

Open-File Report 2020–1053

U.S. Department of the Interior
U.S. Geological Survey

Adjusted Geomagnetic Data—Theoretical Basis and Validation

By E. Joshua Rigler and Abram E. Claycomb

Geomagnetism Program

Open-File Report 2020–1053

**U.S. Department of the Interior
U.S. Geological Survey**

U.S. Department of the Interior
DAVID BERNHARDT, Secretary

U.S. Geological Survey
James F. Reilly II, Director

U.S. Geological Survey, Reston, Virginia: 2020

For more information on the USGS—the Federal source for science about the Earth, its natural and living resources, natural hazards, and the environment—visit <https://www.usgs.gov> or call 1–888–ASK–USGS.

For an overview of USGS information products, including maps, imagery, and publications, visit <https://store.usgs.gov/>.

Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

Although this information product, for the most part, is in the public domain, it also may contain copyrighted materials as noted in the text. Permission to reproduce copyrighted items must be secured from the copyright owner.

Suggested citation:

Rigler, E.J., and Claycomb, A.E., 2020, Adjusted geomagnetic data—Theoretical basis and validation: U.S. Geological Survey Open-File Report 2020–1053, 19 p., <https://doi.org/10.3133/ofr20201053>.

ISSN 2331-1258 (online)

Contents

Abstract.....	1
Motivation	1
Traditional Baseline Adjustments	2
Affine Transformations.....	5
Estimating Affine Transformation.....	7
Ordinary Least Squares	7
Singular Value Decomposition	9
Adaptive Affine Matrices.....	11
Weighted Least Squares	11
Affine Matrix Interpolation.....	11
Adjusting Data	12
Synthetic Data Demonstration	12
Real Data Validation	15
Summary and Conclusions.....	18
Acknowledgments.....	19
References Cited.....	19

Figures

1. Graph showing horizontal geomagnetic vectors in variometer and absolute coordinate frames, with traditional baseline adjustments H_{base} and D_{base}	3
2. Graph showing assumed horizontal geomagnetic vectors in variometer and absolute coordinate frames, with traditional baseline adjustments H_{base} and D_{base}	4
3. Graph showing assumed horizontal geomagnetic vectors in variometer and absolute coordinate frames, with traditional baseline adjustments H_{base} and D_{base}	5
4. Graphs showing synthetic “truth” vector components plotted in two-dimensional planar views	12
5. Graphs showing transformed synthetic “truth” or “observed” vector components in two-dimensional planar views.....	13
6. Graph showing synthetic “truth” and “observed” vector components plotted in three dimensions	14
7. Graphs of sampled data showing the “truth” series at irregular times (similar to absolute measurements), and the “observed” series at those same times (similar to variometer measurements)	14
8. Graphs showing differences between four raw data correction methods and static traditional baseline adjustments.....	16
9. Graph showing the difference between measured and calculated total field	17

Abbreviations

BOU	geomagnetic observatory near Boulder, Colorado
INTERMAGNET	International Real Time Magnetic Observatory Network
MAD	mean absolute ΔF deviation
nT	nanoTesla
RMS	root-mean-square
Slerp	spherical linear interpolation
USGS	U.S. Geological Survey

Adjusted Geomagnetic Data—Theoretical Basis and Validation

By E. Joshua Rigler and Abram E. Claycomb

Abstract

Adjusted geomagnetic data are magnetometer measurements with provisional correction factors applied such that vector quantities are oriented in a local Cartesian frame in which the X axis points north, the Y axis points east, and the Z axis points down. These correction factors are determined from so-called absolute measurements, which are “ground truth” observations made in the field using specialized magnetometers and survey equipment that are (nearly) colocated with the automated and continuously running magnetic measurement instrumentation. Correction factors can be substantial, up to hundreds of nanoTeslas, depending on the geologic and geomagnetic characteristics of the observatory site. They also tend to evolve over time because of instrument response instability and changing site characteristics. Historically, correction factors were determined offline, up to 1 year or more post-measurement, and applied to raw measurements to produce “Definitive” data for scientific analysis. Growing demand for corrected real-time geomagnetic data to better support space weather operations motivated development of an “Adjusted” geomagnetic data product. Modern computational tools, and some notable practical concerns, dictated a transition to affine transformations in lieu of more traditional baseline corrections, as well as a calibration parameter estimation algorithm that is more robust and statistically optimal, and therefore better suited for automated and unsupervised execution. A theoretical basis for this algorithm is presented, along with a demonstration and validation based on a comparison of results obtained with traditional techniques. Discrepancies between Definitive corrected data and near real-time Adjusted data obtained using affine transformations are minimal, generally much less than 5 nanoTeslas per vector component, and less than 1 nanoTesla for the total field magnitude, which satisfies International Real-Time Magnetic Observatory Network (INTERMAGNET) standards.

Motivation

The U.S. Geological Survey (USGS) Geomagnetism Program employs 3-axis vector fluxgate magnetometers at 14 sites throughout the United States and U.S. territories (Love and others, 2008; Love and Finn, 2011). Alone, these magnetometers are high-quality variometers. That is, they record the relative variation of Earth's magnetic field over time very precisely, but are notoriously difficult to calibrate in an absolute sense. Only when variometer data are combined with frequent absolute calibration measurements are geomagnetic observatory quality data possible (Jankowsky and Sucksdorff, 1996).

Historically, the merging of these absolute calibrations and magnetic variation measurements has been a laborious process, ultimately resulting in the magnetic observatory community's standard “Definitive” data product for each observatory. Definitive data are processed in 1-year blocks, starting no less than 1 year after the first observation made that year, and therefore not nearly a real-time product. More recently, some of the most stringent requirements for Definitive data were relaxed so that a new community standard, “Quasi-Definitive” data, could be produced in a timelier manner (with an approximately 1-month delay), but this is still not real-time enough for many modern technological applications.

The demand for calibrated near real-time geomagnetic data is growing and largely unmet. Following the terminology of the International Real-Time Magnetic Observatory Network (INTERMAGNET), a consortium of international magnetic observatory operators that helps establish operational standards terminology, we refer to this as “Adjusted” data, although this level of processing is also referred to as “Provisional” in the magnetic observatory community. The current INTERMAGNET Technical Reference Manual (St-Louis, 2012) defines Adjusted data as data that have had preliminary calibrations applied in near real-time, where “near real-time” means fewer than 7 days. A nearly completed updated edition of the INTERMAGNET manual includes a similar description, but encourages a more stringent definition of near real-time of less than 2 minutes, although the draft manual acknowledges that few observatories meet this goal and INTERMAGNET itself cannot disseminate the data nearly so quickly. The draft manual also encourages use of Quasi-Definitive data if such data can be produced in near

2 Adjusted Geomagnetic Data—Theoretical Basis and Validation

real-time. We have chosen to use this latter criterion to gauge success in this study. That is, our Adjusted data should satisfy the statistical specifications of Quasi-Definitive data, but we acknowledge and accept that real-time data will be less accurate and (or) have more gaps.

This report is an attempt to (1) distill and document institutional knowledge associated with legacy procedures and software for generating Definitive geomagnetic data; (2) assess how or why these legacy procedures and software are inadequate for near real-time Adjusted data generation; and (3) describe, demonstrate, and validate a new methodology that is more tailored for near real-time operations and data processing.

Traditional Baseline Adjustments

Fluxgate sensors have a limited range over which their response functions can be considered linear. In order to maximize their sensitivity while still measuring the extremely wide range of magnetic environments encountered across Earth's surface, common practice is to generate a bias current and magnetic field that opposes and (mostly) nullifies the Earth's main magnetic field along each of the instrument's axes. In doing so, the fluxgate can precisely capture minute relative variations in the magnetic field relative to the main field. One need only add the static oppositional field back to these variational measurements for an accurate and precise measure of the time-varying geomagnetic field.

Although this oppositional magnetic field strength might be estimated from first principals, significant static and time-dependent uncertainties in the fluxgate's mechanical geometries, plus electronic/electrical inefficiencies, and even instability of the sensor's mounting platform, conspire to require frequent and regular calibration to meet tolerances expected of magnetic observatories. So frequent, in fact, that removing the sensors to a laboratory for calibration is not generally advisable, so the sensors are effectively calibrated in place using absolute measurements, along with various physical and geometric assumptions, to calculate so-called baseline adjustments.

Absolute measurements include an estimation of the three-dimensional orientation of the magnetic field at a particular time and place, combined with an independent measurement of the total field. This total field measurement is made with an instrument that is generally more stable than a fluxgate, and which is effectively calibrated to a laboratory standard by fundamental physical principles. Often this is a proton-precession type instrument in which the frequency of precession leads to an oscillating voltage. The frequency of this oscillation uniquely determines the total field magnitude, and the only source of error is in the timing electronics. Orientation and the total field magnitude are combined into a vector measurement to which the fluxgate measurements are corrected.

Traditionally, manual absolute measurements should not be made too closely to the total field sensor, or magnetic contamination is a risk. Thus, part of a baseline adjustment implicitly includes a quasi-time-stationary vector difference between the total field and absolutes pier locations. Although this value could be obtained through careful vector measurements at both locations, this is not strictly necessary unless one is interested in isolating the fluxgate's true response from local variations in the ambient magnetic field. In general, we are not.

If the initial alignment of the fluxgate and absolute measurement coordinate frames is close, the remaining misalignment can be treated as a linear perturbation and corrected using simple baseline adjustments. When coordinate frames are not well aligned, an initial rotation may be applied to obtain nominal alignment, and simple baseline corrections can be made after that. To this end, the USGS has historically installed fluxgate sensors so that the primary horizontal axis, H , aligns with the local magnetic meridian. This means that the secondary horizontal axis, E (for eastward), need not nullify the Earth's main field in this direction, because, on average, it will be zero, and does not generally exceed the sensor's linear response range. If a zero horizontal gradient is also assumed—therefore, no gradient in the E direction—then no E baseline correction is needed.

Thus, because E is assumed to be a reliable measurement and the assumption that absolute horizontal intensity, H_{abs} , does not differ between the variometer and absolutes pier where measured, finding the angle between the variometer's H axis and absolute north (that is, the “declination baseline”) is possible:

$$D_{base} = D_{abs} - \sin^{-1}\left(\frac{E}{H_{abs}}\right), \quad (1)$$

where

- D_{base} is the estimated baseline declination angle from geographic north,
- D_{abs} is the measured absolute declination angle from geographic north,
- E is the vector component aligned with fluxgate's “eastward” axis, and
- H_{abs} is measured absolute total horizontal intensity.

Likewise, baseline correction H_{base} is estimated by using the Pythagorean theorem to project absolute horizontal intensity H_{abs} onto the fluxgate's H axis and subtracting the fluxgate's horizontal magnetic field:

$$H_{base} = \sqrt{H_{abs}^2 - E^2} - H, \tag{2}$$

where

- H_{base} is the estimated baseline correction along fluxgate's H axis, and
- H is the vector component aligned with fluxgate's "northward" axis.

Equations 1 and 2 are perhaps better illustrated in figure 1, where the blue vector represents the absolute total horizontal magnetic vector in geographic coordinates, the red vector represents the horizontal vector measured by the fluxgate in its own Cartesian coordinates, the green items comprise baseline corrections, and the brown angle D is the declination relative to the fluxgate's H axis or the difference between D_{abs} and D_{base} .

Finally, the third axis, Z , points downward to complete a right-handed coordinate system. The downward absolute magnetic vector component, Z_{abs} , is assumed to align perfectly with the downward fluxgate axis, Z , so a simple baseline correction is all that is needed:

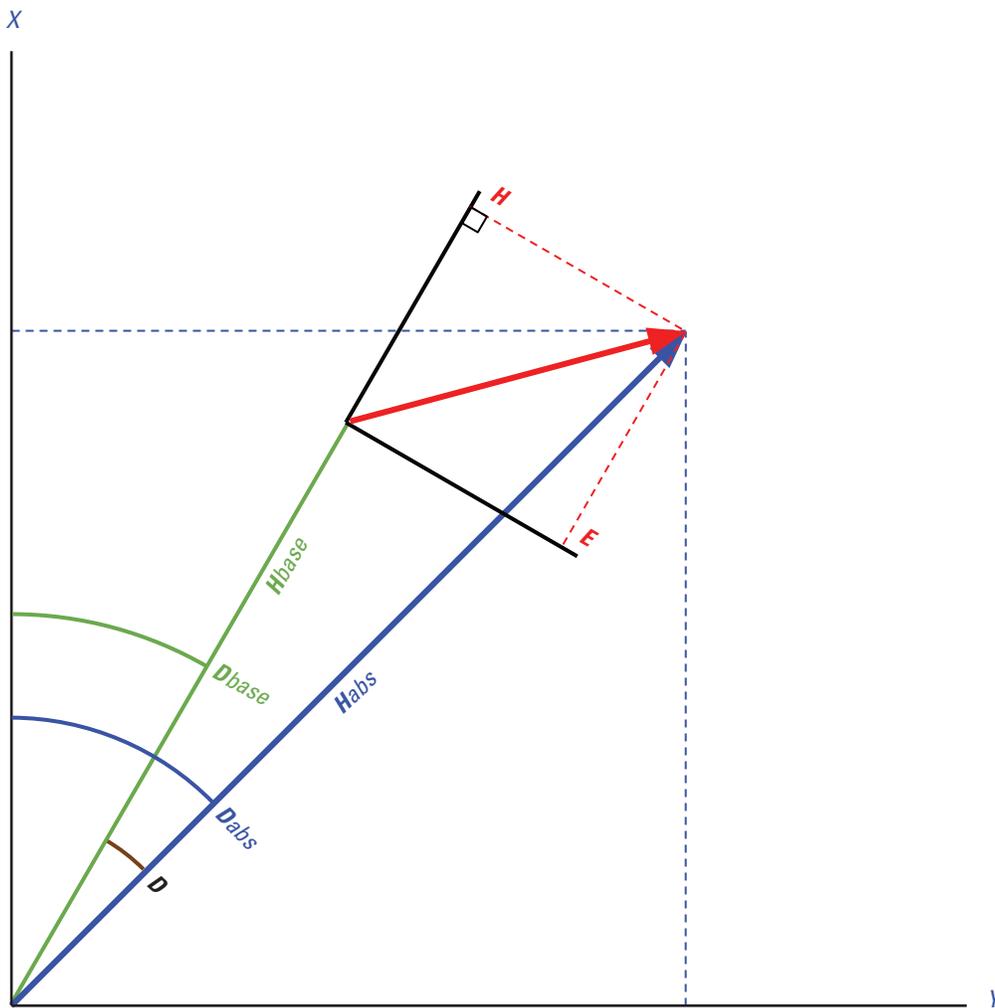


Figure 1. Horizontal geomagnetic vectors in variometer (red) and absolute (blue) coordinate frames, with traditional baseline adjustments H_{base} and D_{base} (green). The blue vector represents the absolute total horizontal magnetic vector in geographic coordinates, the red vector represents the horizontal vector measured by the fluxgate in its own Cartesian coordinates, the green items comprise baseline corrections, and the brown angle D is the declination relative to the fluxgate's H axis or the difference between D_{abs} and D_{base} (D_{abs} , absolute declination angle; D_{base} , declination angle baseline; D , declination difference; E , fluxgate "eastward" axis; H_{abs} , absolute horizontal intensity; H_{base} , horizontal intensity baseline; H , fluxgate "northward" axis)

4 Adjusted Geomagnetic Data—Theoretical Basis and Validation

$$Z_{base} = Z_{abs} - Z, \quad (3)$$

where

- Z_{base} is the estimated “downward” intensity baseline,
- Z_{abs} is the measured absolute downward intensity, and
- Z is the fluxgate’s “downward” vector component.

The assumptions and simplifications invoked in the preceding description of traditional geomagnetic baseline corrections are not always valid. For example, what if an unknown bias current has been applied to the fluxgate’s E axis? This would lead to an incorrect projection of H_{abs} onto the fluxgate’s H axis and therefore an invalid H_{base} correction. Perhaps more problematic, what if horizontal gradients between the variometer and absolutes pier are not zero? Even if one assumes no baseline correction is needed for the fluxgate’s E axis, if H_{abs} is not the same at the variometer and absolutes pier, then both H_{base} and D_{base} will be invalid. These issues are illustrated, in an admittedly exaggerated manner, in figures 2 and 3.

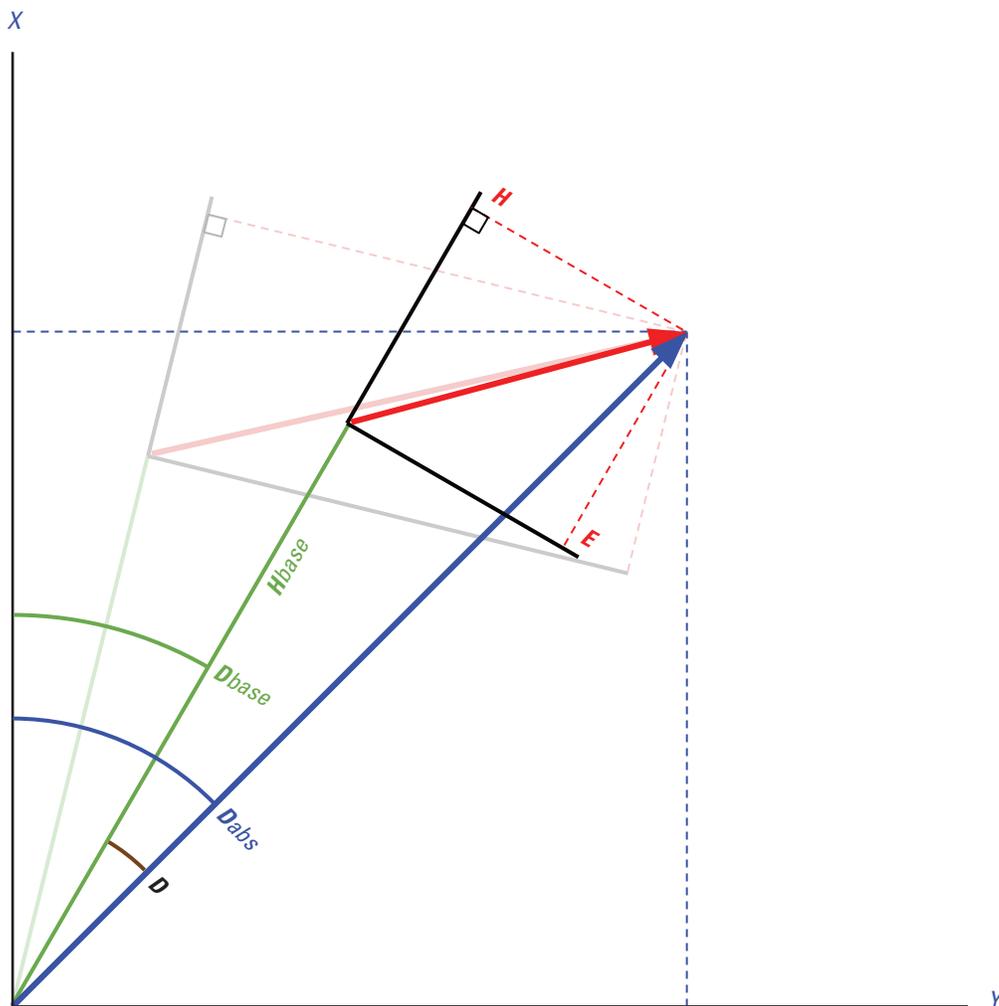


Figure 2. Assumed horizontal geomagnetic vectors in variometer (red) and absolute (blue) coordinate frames, with traditional baseline adjustments H_{base} and D_{base} (green). Actual vectors and frames are shown in faded colors when an unknown bias current corrupts the fluxgate’s E axis. The blue vector represents the absolute total horizontal magnetic vector in geographic coordinates, the red vector represents the horizontal vector measured by the fluxgate in its own Cartesian coordinates, the green items comprise baseline corrections, and the brown angle D is the declination relative to the fluxgate’s H axis or the difference between D_{abs} and D_{base} .

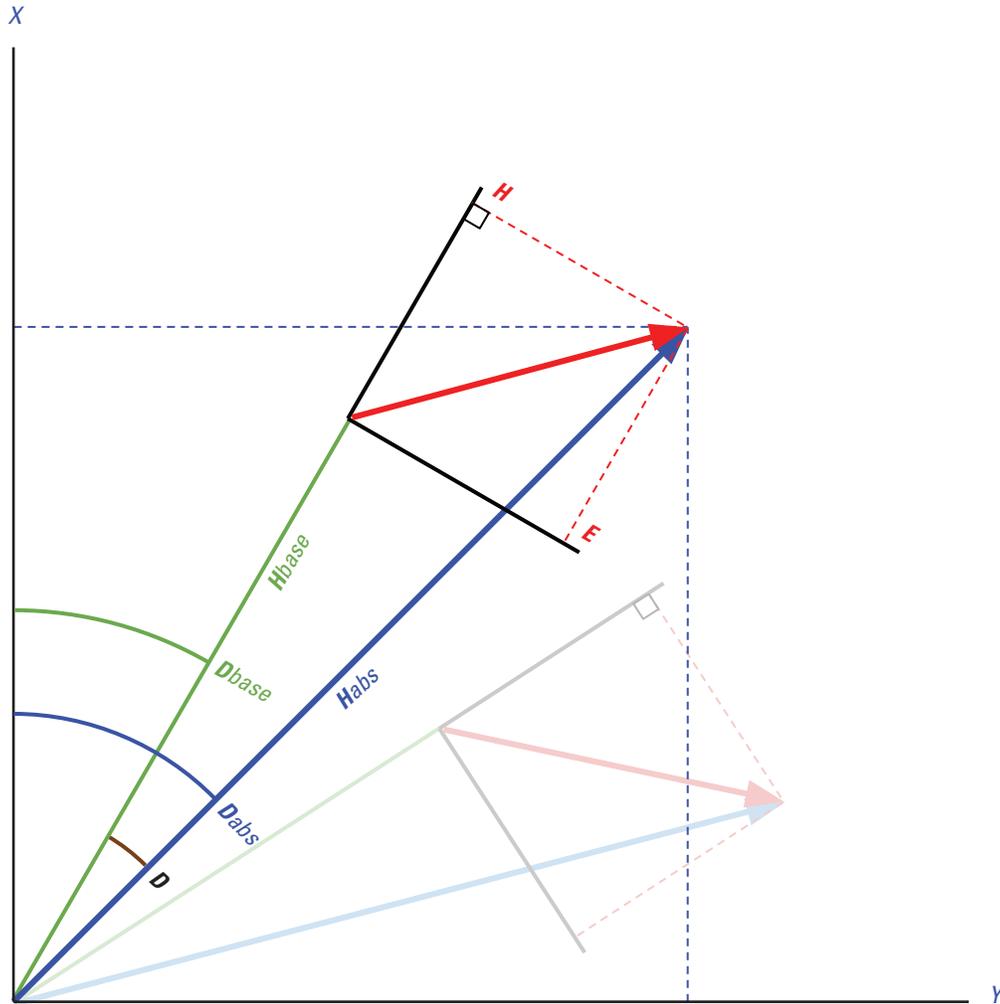


Figure 3. Assumed horizontal geomagnetic vectors in variometer (red) and absolute (blue) coordinate frames, with traditional baseline adjustments H_{base} and D_{base} (green). Actual vectors and frames are shown in faded colors when a gradient is in the horizontal ambient geomagnetic field. The blue vector represents the absolute total horizontal magnetic vector in geographic coordinates, the red vector represents the horizontal vector measured by the fluxgate in its own Cartesian coordinates, the green items comprise baseline corrections, and the brown angle D is the declination relative to the fluxgate's H axis or the difference between D_{abs} and D_{base} .

Finally, for completeness, we acknowledge that not uncommonly small angle approximations are invoked when calculating H_{base} or D_{base} , because, for most observatories, the ratio of the E to the H_{abs} magnetic vector component is small. For example, even for the observatory at Barrow, Alaska, small angle error is rarely more than a few parts per million when the fluxgate is aligned with the local magnetic meridian. However, the local magnetic meridian drifts over time, more so at higher latitudes than at lower latitudes. Furthermore, as one approaches the magnetic pole, the concept of a magnetic meridian becomes undefined. If traditional baseline correction procedures are followed, best practice is to avoid small angle approximations altogether.

Affine Transformations

Even when the assumptions stated previously (that is, no bias current offset in E and zero horizontal gradient) are valid, other sources of uncertainty can complicate traditional baseline adjustments. For example, typically each fluxgate axis is assumed to be perfectly independent of the other two, but in reality, sensor axes are rarely perfectly orthogonal. Also, the scale

6 Adjusted Geomagnetic Data—Theoretical Basis and Validation

factors that convert actual measured voltages for each axis into physical units may not be perfectly calibrated. In fact, if one considers separately the non-orthogonality of the axes, scaling differences for each axis, coordinate frame rotation, and actual baseline differences between the fluxgate sensor and the absolute measurement device, no fewer than 12 degrees of freedom might be considered.

A more rigorous approach to adjusting raw magnetic vector data is to generate a linear transformation matrix \mathbf{M} that directly converts variometer data from the data's own local Cartesian HEZ sensor coordinates into absolutely calibrated Cartesian XYZ geographic coordinates. Although not common, the magnetic observatory community has precedent for taking this approach (Hitchman and others, 2013). Standard linear transformations involve rotation, scaling, and even shear (that is, nonrigid rotation), and alignment of the two coordinate frame origins (that is, translation) can be easily included with a simple augmentation of the measurement vectors and \mathbf{M} . This is known as an affine transformation:

$$\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{M} \begin{bmatrix} H \\ E \\ Z \\ 1 \end{bmatrix}, \quad (4)$$

where

- \mathbf{M} is a composition of potentially many separate affine transformations;
- HEZ is a vector triple of measurements aligned with the fluxgate's H , E , and Z axes; and
- XYZ is a vector triple of measurements aligned with absolute geographic coordinate X , Y , and Z axes.

$\mathbf{M} =$

$$\begin{bmatrix} 1 & 0 & 0 & t_X \\ 0 & 1 & 0 & t_Y \\ 0 & 0 & 1 & t_Z \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot (i.e., \mathbf{T}) \cdot \left[\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta_X & -\sin\theta_X & 0 \\ 0 & \sin\theta_X & \cos\theta_X & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \cos\theta_Y & 0 & \sin\theta_Y & 0 \\ 0 & 1 & 0 & 0 \\ -\sin\theta_Y & 0 & \cos\theta_Y & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \cos\theta_Z & -\sin\theta_Z & 0 & 0 \\ \sin\theta_Z & \cos\theta_Z & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot (i.e., \mathbf{R}) \right] \cdot \begin{bmatrix} s_X & 0 & 0 & 0 \\ 0 & s_Y & 0 & 0 \\ 0 & 0 & s_Z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot (i.e., \mathbf{S}) \cdot \begin{bmatrix} 1 & h_{XY} & h_{XZ} & 0 \\ 0 & 1 & h_{YZ} & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} (i.e., \mathbf{H}), \quad (5)$$

where

- \mathbf{T} is a translation affine transformation matrix, and
- t_X is the translation along the X axis;
- \mathbf{R} is a rotation affine transformation matrix, and
- θ_X is a rotation angle around the X axis;
- \mathbf{S} is a scaling affine transformation matrix, and
- s_X is a scale factor for the X axis;
- \mathbf{H} is a shear affine transformation matrix, and
- h_{XY} is a shear in the XY plane, aligned with the X axis.

Similar definitions apply to the remaining variables.

Note that the order in which the constituent matrices are multiplied when composing an affine transformation is critical. However, no single standard exists for composing an affine transformation from its constituent matrices. Equation 5 follows a norm used in the Python Transforms3D package (Brett, 2016) and might better be represented algorithmically as $\text{dot}(\mathbf{T}, \text{dot}(\mathbf{R}, \text{dot}(\mathbf{S}, \mathbf{H})))$, where “dot” is a matrix dot product function. In words, a shear correction (\mathbf{H}) is first applied to the original HEZ

vector to orthogonalize it. This is followed by rescaling the orthogonalized HEZ vector with \mathbf{S} , which is followed by a rigid rotation \mathbf{R} to align all axes with the desired XYZ coordinate frame. Finally, the origin of the now rotated, scaled, and orthogonal fluxgate vector is translated to make it coincide with the origin of the absolute measurement frame using \mathbf{T} .

Looking more closely at the rotation matrix $\mathbf{R} = \text{dot}(\mathbf{R}_X, \text{dot}(\mathbf{R}_Y, \mathbf{R}_Z))$, it should be even more evident that the order in which the rotations are applied is important. The rotation angles (θ_X , θ_Y , and θ_Z) that combine to form \mathbf{R} are presented here in a manner that resembles typical yaw-pitch-roll rotations. This is one of six possible so-called Tait-Bryan rotation sequences (reorientation by rotating around each unique axis once). Proper Euler angles have six additional possible sequences that involve a rotation about one axis, then a second, then again about the first. If that were not enough, both the Tait-Bryan and proper Euler angles can be with respect to fixed axes (extrinsic, as the previous example), or with respect to the new axes after each rotation (intrinsic). This leads to 24 possible valid rotations! We recommend that one convention be chosen and used consistently.

The reader may wonder at the particular form of \mathbf{H} . Allowing shear coefficients in all the off-diagonal elements of the upper-left 3x3 matrix is certainly possible, but we chose not to here for at least two reasons: (1) shear, by definition, must not alter the volume of the original data points, or in mathematical terms, its determinant is always 1, which is guaranteed by this triangular form; and (2) this triangular form permits an efficient mechanism for separating rigid rotation from pure shear (Thomas, 1991). A nuance related to the second reason is that the solution is unique only insofar as this particular matrix structure is chosen to represent shear. If a different structure were chosen, the rotation matrix will be different, even though the composition of shear and rotation will always be the same. In more practical terms, this means that speaking of “rotation” and “shear” separately is not correct, and dropping one without careful consideration is not appropriate. In particular, the structure chosen here (that is, upper triangular) means that the vector is first rotated so that the Z axes of both coordinate frames are aligned perfectly. Then, the E axis is sheared in the Z direction, and the H axis is sheared in both Z and E directions.

Estimating Affine Transformation

The constituent affine matrices can, in theory, be created individually, through careful sensor design and construction, laboratory calibration, and (or) surveying of the observatory site, but in practice these actions are not always possible or adequate. Ultimately, if absolute XYZ measurements are considered the “truth” to which variational measurements are to be corrected, the linear structure of the affine transformation allows us to invoke linear estimation theory to determine an optimal transformation matrix \mathbf{M} that encapsulates all the necessary corrections to the original HEZ vector data:

$$\begin{bmatrix} X_1 & X_2 & \cdots & X_n \\ Y_1 & Y_2 & \cdots & Y_n \\ Z_1 & Z_2 & \cdots & Z_n \\ 1 & 1 & \cdots & 1 \end{bmatrix} = \begin{bmatrix} M_{11} & M_{12} & M_{13} & M_{14} \\ M_{21} & M_{22} & M_{23} & M_{24} \\ M_{31} & M_{32} & M_{33} & M_{34} \\ M_{41} & M_{42} & M_{43} & M_{44} \end{bmatrix}^x \begin{bmatrix} H_1 & H_2 & \cdots & H_n \\ E_1 & E_2 & \cdots & E_n \\ Z_1 & Z_2 & \cdots & Z_n \\ 1 & 1 & \cdots & 1 \end{bmatrix}, \quad (6)$$

where

n is the length of the HEZ and XYZ series.

Ordinary Least Squares

The most straightforward estimation technique is least squares. Many mathematical software libraries contain efficient routines that, given measurement arrays constructed similarly to equation 6, will use linear least squares to determine an optimal 4x4 \mathbf{M} matrix to map $HEZ1$ into $XYZ1$ vectors. However, what all these routines inevitably do, at least “under the hood,” is rearrange the previous matrix equation algebraically such that:

$$\begin{bmatrix} X_1 \\ Y_1 \\ Z_1 \\ 1 \\ X_2 \\ Y_2 \\ Z_2 \\ 1 \\ \vdots \\ X_n \\ Y_n \\ Z_n \\ 1 \end{bmatrix} = \begin{bmatrix} H_1 & E_1 & Z_1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & H_1 & E_1 & Z_1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_1 & E_1 & Z_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_1 & E_1 & Z_1 \\ H_2 & E_2 & Z_2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & H_2 & E_2 & Z_2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_2 & E_2 & Z_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_2 & E_2 & Z_2 \\ \vdots & \vdots \\ H_n & E_n & Z_n & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & H_n & E_n & Z_n & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_n & E_n & Z_n & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_n & E_n & Z_n \end{bmatrix} \cdot \begin{bmatrix} M_{11} \\ M_{21} \\ M_{31} \\ M_{41} \\ M_{21} \\ M_{22} \\ M_{23} \\ M_{24} \\ M_{31} \\ M_{32} \\ M_{33} \\ M_{34} \\ M_{41} \\ M_{42} \\ M_{43} \\ M_{44} \end{bmatrix}. \quad (7)$$

Equation 7 is equivalent to equation 6 (that is, a 4x4 affine matrix \mathbf{M}), but perhaps makes it clearer that a set of linear equations is being solved for 16 unknowns. Given that a single vector measurement comprises only three actual data points, plus the augmented “1,” it should also be clear that four sets of absolute and fluxgate measurements are required to obtain a solution.

Constraining the system of equations is often desirable, usually to reduce the degrees of freedom and reduce the uncertainty associated with the solution to a given system of equations. In fact, although one of the biggest advantages to using affine transformations might be that they include scale factors and shear corrections, in addition to rotation and translation, allowing so many free parameters leads to overfitting of the limited training data in our experience. When this happens, the solution, \mathbf{M} , does not generalize well when used to adjust raw data between absolute measurements. If the frequency of absolute measurements could be increased substantially, it might reduce uncertainty enough that realistic scale and shear calibrations could be obtained from absolute measurements. Until such time, limiting the number of degrees of freedom helps avoid overfitting the limited absolute training data.

Such constraints can be made somewhat consistent with traditional Definitive data processing. To start, if we look closely at all the affine transformation matrices in the previous subsection, evidently the final row of \mathbf{M} should always be $[0 \ 0 \ 0 \ 1]$. We can exploit this knowledge to remove the M_{4*} unknowns from our system of equations, and their corresponding columns in the measurement matrices, thus reducing our number of unknowns from 16 to 12, and slightly reducing uncertainty in the remaining estimated coefficients:

$$\begin{bmatrix} X_1 \\ Y_1 \\ Z_1 \\ X_2 \\ Y_2 \\ Z_2 \\ \vdots \\ X_n \\ Y_n \\ Z_n \end{bmatrix} = \begin{bmatrix} H_1 & E_1 & Z_1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & H_1 & E_1 & Z_1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_1 & E_1 & Z_1 & 1 \\ H_2 & E_2 & Z_2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & H_2 & E_2 & Z_2 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_2 & E_2 & Z_2 & 1 \\ \vdots & \vdots \\ H_n & E_n & Z_n & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & H_n & E_n & Z_n & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & H_n & E_n & Z_n & 1 \end{bmatrix} \cdot \begin{bmatrix} M_{11} \\ M_{21} \\ M_{31} \\ M_{41} \\ M_{21} \\ M_{22} \\ M_{23} \\ M_{24} \\ M_{31} \\ M_{32} \\ M_{33} \\ M_{34} \end{bmatrix}. \quad (8)$$

This simple constraint does not buy us much. But now suppose we wish to reduce the number of free parameters even further by imposing additional prior knowledge/assumptions on the structure of our matrix that are (mostly) consistent with traditional Definitive data processing techniques. For example, assume that the vertical vector component Z is perfectly aligned between the variometer and the absolute instrument, and it is perfectly orthogonal to the fluxgate’s horizontal components. In

addition, assume that the transformation of the horizontal vector component is simply a scaled rotation about the Z axis (that is, no shear). Finally, allow translation of the vertical component, but don't allow any translation of the horizontal components. This is accomplished by removing and (or) moving elements of equation 8 to give:

$$\begin{bmatrix} X_1 \\ Y_1 \\ (Z-Z)_1 \\ X_2 \\ Y_2 \\ (Z-Z)_2 \\ \vdots \\ X_n \\ Y_n \\ (Z-Z)_n \end{bmatrix} = \begin{bmatrix} H_1 & E_1 & \cdots & 0 \\ E_1 & -H_1 & \cdots & 0 \\ 0 & 0 & \cdots & 1 \\ H_2 & E_2 & \cdots & 0 \\ E_2 & -H_2 & \cdots & 0 \\ 0 & 0 & \cdots & 1 \\ \vdots & \vdots & \cdots & \vdots \\ H_n & E_n & \cdots & 0 \\ E_n & -H_n & \cdots & 0 \\ 0 & 0 & \cdots & 1 \end{bmatrix} \begin{bmatrix} M_{11} \\ M_{21} \\ \vdots \\ \vdots \\ M_{34} \end{bmatrix}. \quad (9)$$

In equation 9, green indicates moved/edited elements, and red indicates whole columns of inputs that were removed, which corresponds to removal of one of the M coefficients. Also note that the Z component from the variometer is subtracted from the Z component obtained from absolute measurements to provide a Z translation that is completely independent of the X and Y axes. The degrees of freedom have been reduced from 16 to just 3.

Once data matrices are arranged properly, ordinary least squares is invoked to solve for the M coefficients. This is a commonly used method, available in almost all mathematical software libraries. We present it here in matrix form for completeness:

$$\hat{M} = (X^T X)^{-1} X^T Y, \quad (10)$$

where

- \hat{M} is the solution vector that minimizes the sum of squared errors,
- X is the independent variable matrix containing the HEZ variometer observations, and
- Y is the dependent variable vector containing the XYZ absolute observations.

Once \hat{M} is obtained, its coefficients should be inserted into their corresponding locations in the two-dimensional \hat{M} affine matrix:

$$\hat{M} = \begin{bmatrix} M_{11} & M_{12} & 0 & 0 \\ -M_{12} & M_{11} & 0 & 0 \\ 0 & 0 & 1 & M_{34} \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (11)$$

Singular Value Decomposition

We emphasize now that solving this reduced set of equations is not exactly equivalent to traditional Definitive data processing as presented earlier. In fact, such an affine transformation would look something like:

$$M = \begin{bmatrix} \cos \theta_Z & -\sin \theta_Z & 0 & t_X \\ \sin \theta_Z & \cos \theta_Z & 0 & t_Y \\ 0 & 0 & 1 & t_Z \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (12)$$

In equation 12, t_X and t_Y are not independent, but are geometrically linked to each other through θ_Z (that is, as θ_Z approaches 0, t_Y approaches 0 and t_X approaches H_{base}). But now equation 12 cannot be solved directly from the measurement matrices using simple linear least squares because it is no longer linear in its parameters. However, θ_Z can be found by invoking singular value decomposition to obtain the eigenvectors that define an orthonormal rotation matrix between two vector spaces (for example, Eggert and others, 1997).

10 Adjusted Geomagnetic Data—Theoretical Basis and Validation

With modern numerical libraries, this is relatively simple. First, remove the means from both the HE and XY measurement matrices so that their origins coincide (ignore Z for now):

$$\begin{aligned} H'_i &= H_i - \bar{H}, & E'_i &= E_i - \bar{E} \\ X'_i &= X_i - \bar{X}, & Y'_i &= Y_i - \bar{Y} \end{aligned} \quad (13)$$

where

H'_i is a series of H measurements with its mean removed,
 H_i is a series of H measurements, and
 \bar{H} is the mean of series H_i .

Similar definitions apply to the remaining variables.

Next, create a cross-covariance matrix between HE and XY :

$$\mathbf{C} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{bmatrix}, \quad (14)$$

where

\mathbf{C} is the cross-covariance matrix between HE and XY ; and

$$c_{11} = \sum_i^n H'_i X'_i, \quad c_{12} = \sum_i^n H'_i Y'_i, \quad c_{21} = \sum_i^n E'_i X'_i, \quad c_{22} = \sum_i^n E'_i Y'_i, \quad (15)$$

where

c_{11} is the cross-covariance between H and X ,
 c_{12} is the cross-covariance between H and Y ,
 c_{21} is the cross-covariance between E and X , and
 c_{22} is the cross-covariance between E and Y .

Next, decompose \mathbf{C} into its singular values, plus left and right eigenvectors using the singular-value-decomposition routine found in your favorite numerical linear algebra library:

$$\mathbf{C} = \mathbf{U}\mathbf{S}\mathbf{V}^* \quad (16)$$

where

\mathbf{U} is the left eigenvector matrix of \mathbf{C} ,
 \mathbf{S} is the singular value (or eigenvalue) matrix of \mathbf{C} , and
 \mathbf{V}^* is the right eigenvector matrix of \mathbf{C} .

In [equation 16](#), singular values s_i compose the diagonal matrix \mathbf{S} , and the orthonormal matrices \mathbf{U} and \mathbf{V}^* can be combined to give the unscaled rotation that aligns HE with XY :

$$\mathbf{R} = \mathbf{V}^{*\top} \mathbf{U}^\top. \quad (17)$$

Finally, using \mathbf{R} , simply rotate \bar{H} and \bar{E} into absolute coordinates to obtain t_X and t_Y :

$$\begin{bmatrix} t_X \\ t_Y \end{bmatrix} = \mathbf{R} \cdot \begin{bmatrix} \bar{H} \\ \bar{E} \end{bmatrix}. \quad (18)$$

In [equation 18](#), t_Z is simply the difference in the averages of the vertical components, as with traditional processing.

Notice that there are now four free parameters (θ_Z , t_X , t_Y , and t_Z) instead of just the three traditional baseline corrections. This is because t_X and t_Y depend on both \bar{H} and \bar{E} . By increasing the degrees of freedom from 3 to 4, the uncertainty of the final solution is somewhat higher. However, uncertainty should be higher because a potential offset in the variometer's E axis that was previously ignored is now considered in the calculations.

Adaptive Affine Matrices

Even after switching to affine matrices to transform *HEZ* variometer vector observations into *XYZ* Adjusted vectors, long-period biases remain in the residuals that often exceed Quasi-Definitive tolerances. In order to remove these biases, and therefore track the non-time-stationary relationship between the variometer and absolute instrumentation, some kind of time-adaptive algorithm is required. What's more, to support arbitrary epoch times, interpolating between less frequent, but optimal, transformation matrices that we will refer to as “keyframes” (a concept borrowed from video compression) is often more computationally efficient.

Weighted Least Squares

Absolute observations tend not to occur regularly, otherwise a recursive least squares algorithm with a finite “memory” might be the best way to estimate non-time-stationary affine matrices. However, because Absolute observations are also relatively infrequent (sampling intervals on the order of days to weeks), building up data matrices from all, or a large subset of all, available observations is not unreasonable; then estimating the optimal transformation for a given interval using ordinary least squares becomes feasible. To reduce noise and obtain a more reliable solution, typically many more observations than a bare minimum equal to the number of free parameters are required. Obtaining this number of observations may require a sampling time of several months or more, and because ordinary least squares gives the optimal solution for the entire interval, the keyframe affine matrix estimated for any given epoch may not reflect the true (noise-free) transformation for that epoch.

An estimation method that emphasizes more recent over older measurements would be more representative of the desired epoch, while still using multiple observations to reduce noise and obtain a more reliable solution. Weighted least squares, with weights that are a function of the age of the observation relative to the epoch, is just such a method. Many techniques and software libraries are available to estimate a weighted least squares solution, so we give a general matrix-oriented description here before discussing how to choose weights:

$$\hat{M} = (X^T W X)^{-1} X^T W Y, \quad (19)$$

where

- X is the same independent variable matrix containing the *HEZ* variometer observations as in [equation 10](#),
- W is a diagonal matrix whose elements are the weights to be assigned to each observation, and
- Y is the dependent variable array containing the *XYZ* absolute observation described in [equation 10](#).

If W were the identity matrix, I , weighted least squares reduces to ordinary least squares. The solution, \hat{M} , is actually the vector of matrix coefficients described previously, so each element must then be placed into its corresponding location in the two-dimensional M matrix used to actually adjust the *HEZ* variometer vector measurements.

Affine Matrix Interpolation

Linearly interpolating between scalar values that fall on a common axis is straightforward. Almost all mathematical software libraries implement such a function, so we will not describe any algorithms here. However, problems arise with the rotational components of an affine transformation. Indeed, if linear interpolation were performed for each of the matrix elements used to define a rotation, none of the interpolated matrices would actually be a rotation matrix, and strange, often nonphysical, transformations could result.

Fortunately, this problem has a long-established solution that uses spherical linear interpolation (Slerp; Shoemake, 1985). Briefly, the Slerp technique finds the vector axis of rotation that characterizes each given rotation matrix, then rotates from the first to the last at a constant rate as if the axis of rotation were tracing out a unit sphere. Slerp is not nearly as common as standard linear interpolation, but the spatial transformations subpackage in Scientific Python (for example, Oliphant, 2007; Millman and Aivazis, 2011) provides a robust implementation.

Assuming one begins with an arbitrary affine transformation matrix, decomposing the matrix into scale, rotation, and translational components first is necessary. Matrices can be decomposed in several ways, including singular value decomposition, as described previously. A simpler and more robust method involves so-called polar decomposition (Shoemake and Duff, 1992). Once decomposed, standard linear interpolation is performed on the scale and translation matrices, and Slerp is applied to the rotation matrix. The constituent transformations are recombined to give a valid, more physically consistent affine transformation at each intermediate time step.

Adjusting Data

To verify that our theoretical basis is sound and to offer a helpful visual interpretation of how affine transformations work, we present a demonstration of the technique applied to synthetic data. Then, we validate the technique using real geomagnetic data and compare the results with those of traditional baseline corrections.

Synthetic Data Demonstration

First, we define a regular series of times t_i , then construct synthetic “truth” vectors that represent the desired results as functions of t :

$$t_i = \left\{ 0, \frac{\pi}{n}, \frac{2\pi}{n}, \dots, \frac{(n-1)\pi}{n}, \pi \right\}, \quad (20)$$

$$X_i^{\text{truth}} = \frac{(2t_i - \pi)}{\pi}, \quad (21)$$

$$Y_i^{\text{truth}} = \sqrt{1 - (X_i^{\text{truth}})^2} \cos 24 t_i, \text{ and} \quad (22)$$

$$Z_i^{\text{truth}} = \sqrt{1 - (X_i^{\text{truth}})^2} \sin 24 t_i, \quad (23)$$

where

- n is the number of synthetic vector samples,
- t_i is the i^{th} of n time steps,
- X_i^{truth} is the i^{th} of n samples of the “true” X vector component,
- Y_i^{truth} is the i^{th} of n samples of the “true” Y vector component, and
- Z_i^{truth} is the i^{th} of n samples of the “true” Z vector component.

These vector components are plotted in their two-dimensional planar views in [figure 4](#).

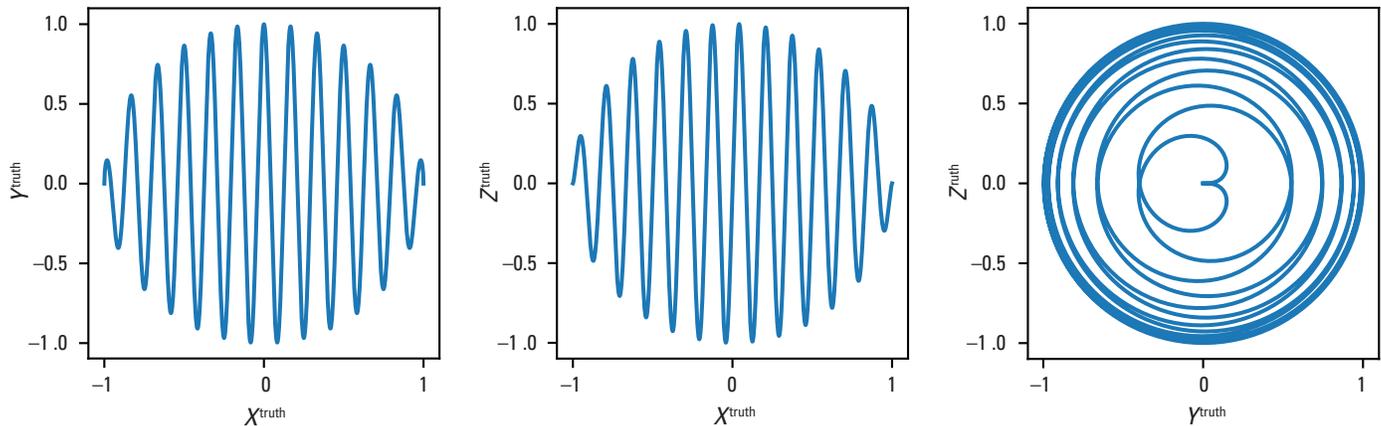


Figure 4. Synthetic “truth” vector components plotted in two-dimensional planar views. (X^{truth} , “true” X vector component; Y^{truth} , “true” Y vector component; Z^{truth} , “true” Z vector component)

Next, we define an affine transformation matrix that converts the truth vectors into a different “observed” basis. This basis is translated from the origin by $\{3,4,3\}$, rotated about an axis defined by the vector $\{1,1,1\}$ by 45 degrees, and scaled by $\{1.5, 0.95, 1.2\}$:

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & 0 & 3 \\ 0 & 1 & 0 & 4 \\ 0 & 0 & 1 & 3 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (24)$$

$$\mathbf{R} = \begin{bmatrix} 0.805 & 0.506 & -0.311 & 0 \\ -0.311 & 0.805 & 0.506 & 0 \\ 0.506 & -0.311 & 0.805 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (25)$$

$$\mathbf{S} = \begin{bmatrix} 1.5 & 0 & 0 & 0 \\ 0 & 0.95 & 0 & 0 \\ 0 & 0 & 1.2 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \text{ and} \quad (26)$$

$$\mathbf{M} = \mathbf{T} \cdot \mathbf{R} \cdot \mathbf{S} = \begin{bmatrix} 1.207 & 0.481 & -0.373 & 3 \\ -0.466 & 0.765 & 0.607 & 4 \\ 0.759 & -0.295 & 0.966 & 3 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (27)$$

Transformed, or “observed,” vector components are plotted in their two-dimensional planar views in [figure 5](#), and a single three-dimensional plot of both synthetic “truth” and “observed” is shown in [figure 6](#).

Finally, the “truth” series is sampled at irregular times (similar to absolute measurements), as is the “observed” series at those same times (similar to variometer measurements). These sampled data are shown in [figure 7](#). These data are used to build the regression matrices, and linear least squares is used to solve for the $\hat{\mathbf{M}}$ coefficients that map the variometer measurements to Adjusted. All coefficients match the true coefficients to several significant digits more than the 3 presented in [equation 27](#). Of course, data will be corrupted by noise in real-world scenarios, and such a perfect match is unlikely, but more observations will reduce uncertainty as long as observation errors are random.

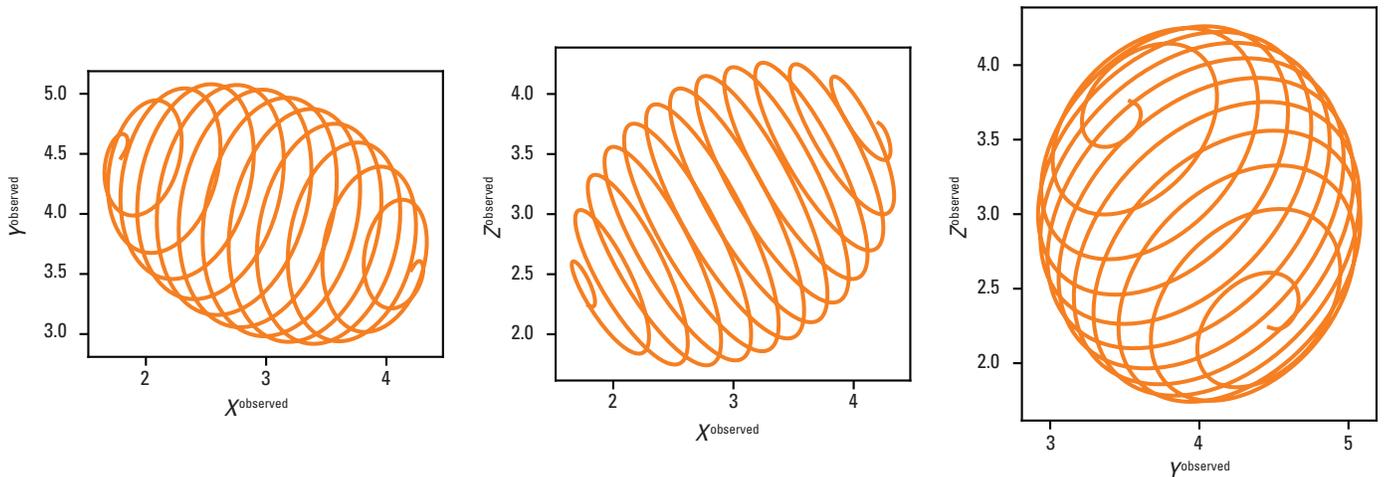


Figure 5. Transformed synthetic “truth” or “observed” vector components in two-dimensional planar views. (X^{observed} , “observed” X vector component; Y^{observed} , “observed” Y vector component; Z^{observed} , “observed” Z vector component)

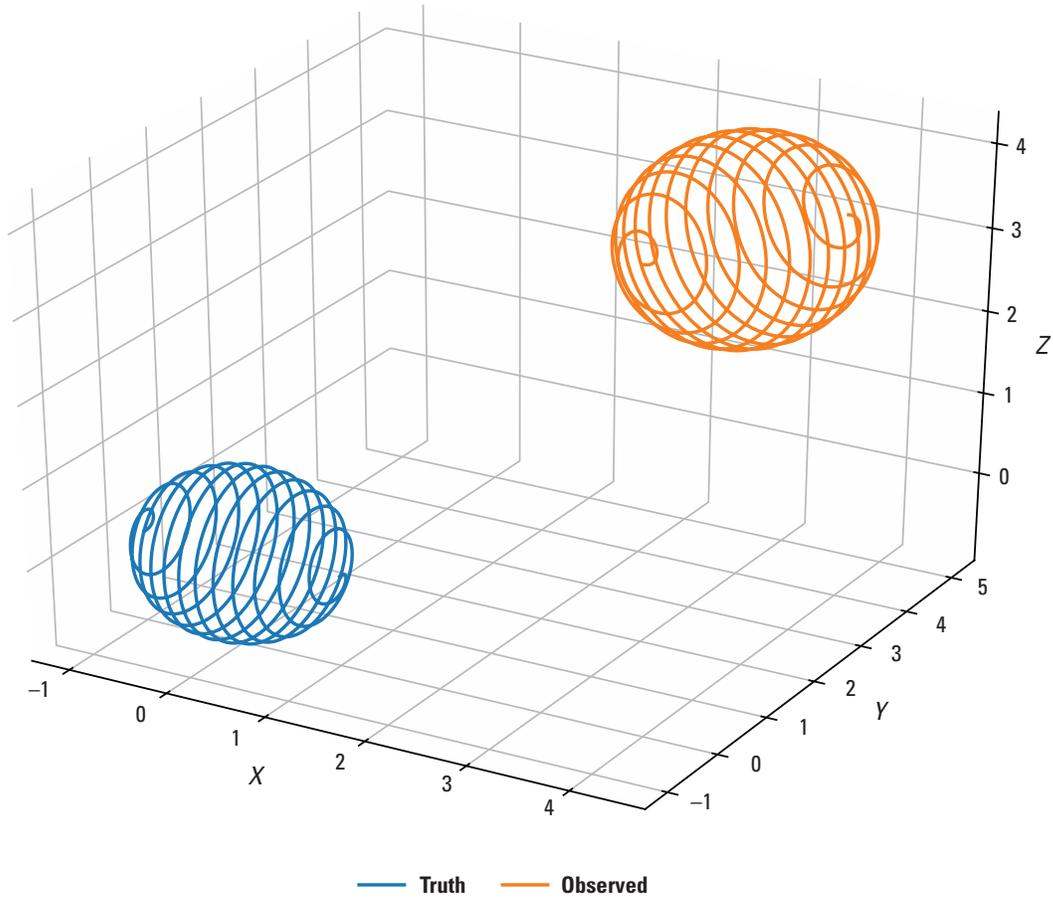


Figure 6. Synthetic “truth” and “observed” vector components plotted in three dimensions.

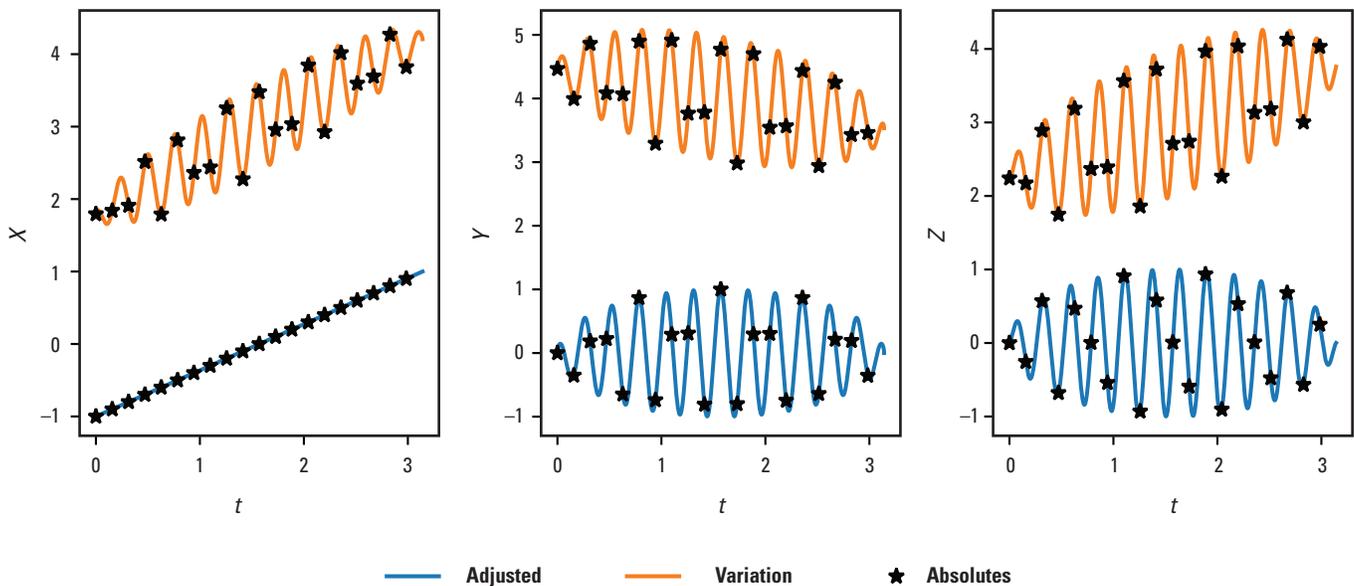


Figure 7. Sampled data showing the “truth” series at irregular times (similar to absolute measurements) and the “observed” series at those same times (similar to variometer measurements). Black stars are samples taken at irregular intervals and are analogous to absolute measurements. Orange traces are synthetic magnetic vector variation observations plotted against time, t . Blue traces are synthetic “truth” and what we expect Adjusted data to resemble in the end. (t , time)

Real Data Validation

An argument has been made for the mathematical rigor and flexibility of using affine transformations to correct raw geomagnetic measurements to match absolute measurements. However, we must acknowledge that this approach is a significant departure from the application of traditional baseline corrections, which have been used successfully for many decades to produce observatory-grade geomagnetic observations. Therefore, using real data to compare the affine transformation method directly with the traditional baseline correction method is important.

For this comparison, we use Definitive data as “ground truth,” even though it may not always be the best representation of reality. Also, this report presents data only for the USGS BOU magnetic observatory near Boulder, Colorado. Similar analyses have been performed for all the USGS observatories; these analyses are available in a Python Notebook that is distributed with versions 0.9.0 or greater of the Geomag-Algorithms software package (<https://github.com/usgs/geomag-algorithms>), as is most of the code used to generate results in this report.

Differences between corrected data are obtained using a static set of traditional baseline adjustments and four alternative adjustment techniques: Definitive, static affine, adaptive affine (acausal), and adaptive affine (causal) (fig. 8). That is, average baseline corrections were estimated for the entire year, then subtracted from the alternative adjustment data series. “Definitive” data are traditional baseline adjustments in which baselines are updated manually to fit each set of absolute measurements; “static affine” data are adjusted with an affine transformation in which the matrix is unchanging and constrained to rigid horizontal rotation plus a translation of origins; “adaptive affine (acausal)” data use an adaptive version of the static affine transformation that automatically fits absolute measurements from both before and after each keyframe epoch, and uses Slerp to interpolate between keyframes (for context, this should be most similar to “Definitive” because it uses absolute measurements from both the past and future); “adaptive affine (causal)” data are adjusted using a version of the static affine transformation that uses absolute measurements from only before each keyframe epoch; nothing is interpolated because no future keyframes exist to which to interpolate. The stars indicate the differences between the static traditional baseline and absolute measurements.

Significant variations in the ΔX and ΔY differences are immediately apparent, especially for the static affine transformation results. This variation also exists in the adaptive affine transformation results, but it is generally smaller, especially for the acausal version. Upon closer inspection, this variation can be seen to be dominated by a daily periodicity. The Definitive results also show a daily variation, but it is usually much less than for any affine transformation results. Why this would be the case is not clear, but it may be at least partly a consequence of the extra degree of freedom mentioned earlier. That is, by including uncertainty in the fluxgate’s E axis, the overall uncertainty of our solution increases slightly. The long-period differences notable in Definitive and the Adaptive affine transformation results are a desired and intended outcome of letting the corrections, either baseline or affine, vary with time.

We present figure 8 largely to illustrate the ability of the different correction techniques to track absolute measurements. The static affine transformation captures the average of the absolute measurements well enough, especially if one considers that most geomagnetic variations at this location are tens to hundreds of nanoTeslas. The causal adaptive affine transformation does a much better job of following the time variable nature of the absolute measurements, although one can clearly see a stair-step characteristic in the fits, as well as a lag relative to changes in the absolute measurements that is a consequence of not including future measurement in the fits. The acausal adaptive affine transformation does an even better job of following absolute measurements: no “stair-steppiness” and no lag.

The Definitive data do appear to follow the absolute measurements most closely, which should not be surprising given that they are the result of many hours of expert analysis. However, a few large deviations in the ΔY panel of figure 8 are difficult to justify physically. A careful review of the analysis used to obtain these data, including a discussion with the responsible analyst, suggests that these deviations may have been the result of poor absolute declination measurements that were not identified and removed in the original analysis. These same measurements were used for all the affine transformation techniques; however, their statistical estimation procedures prevented the solution from being pulled too strongly in a direction that, on review, is clearly nonphysical.

A more traditional measure of overall quality of fit of corrected geomagnetic data, ΔF , is presented in figure 9:

$$\Delta F = \sqrt{X^2 + Y^2 + Z^2} - F, \quad (28)$$

where

ΔF is the difference between the XYZ vector amplitude and F , and
 F is the measured scalar value F .

Judging by the ΔF metric, the affine transformation-corrected data easily satisfy the INTERMAGNET 5-nanoTesla (nT) tolerance. The root mean square (RMS) of the Definitive ΔF for this interval is lowest at 0.29 nT. These same metrics for the static, acausal adaptive, and causal adaptive affine transformations results are 5.15, 4.77, and 4.66 nT, respectively. These might be considered marginal according to INTERMAGNET standards, and therefore give us pause when choosing to employ these

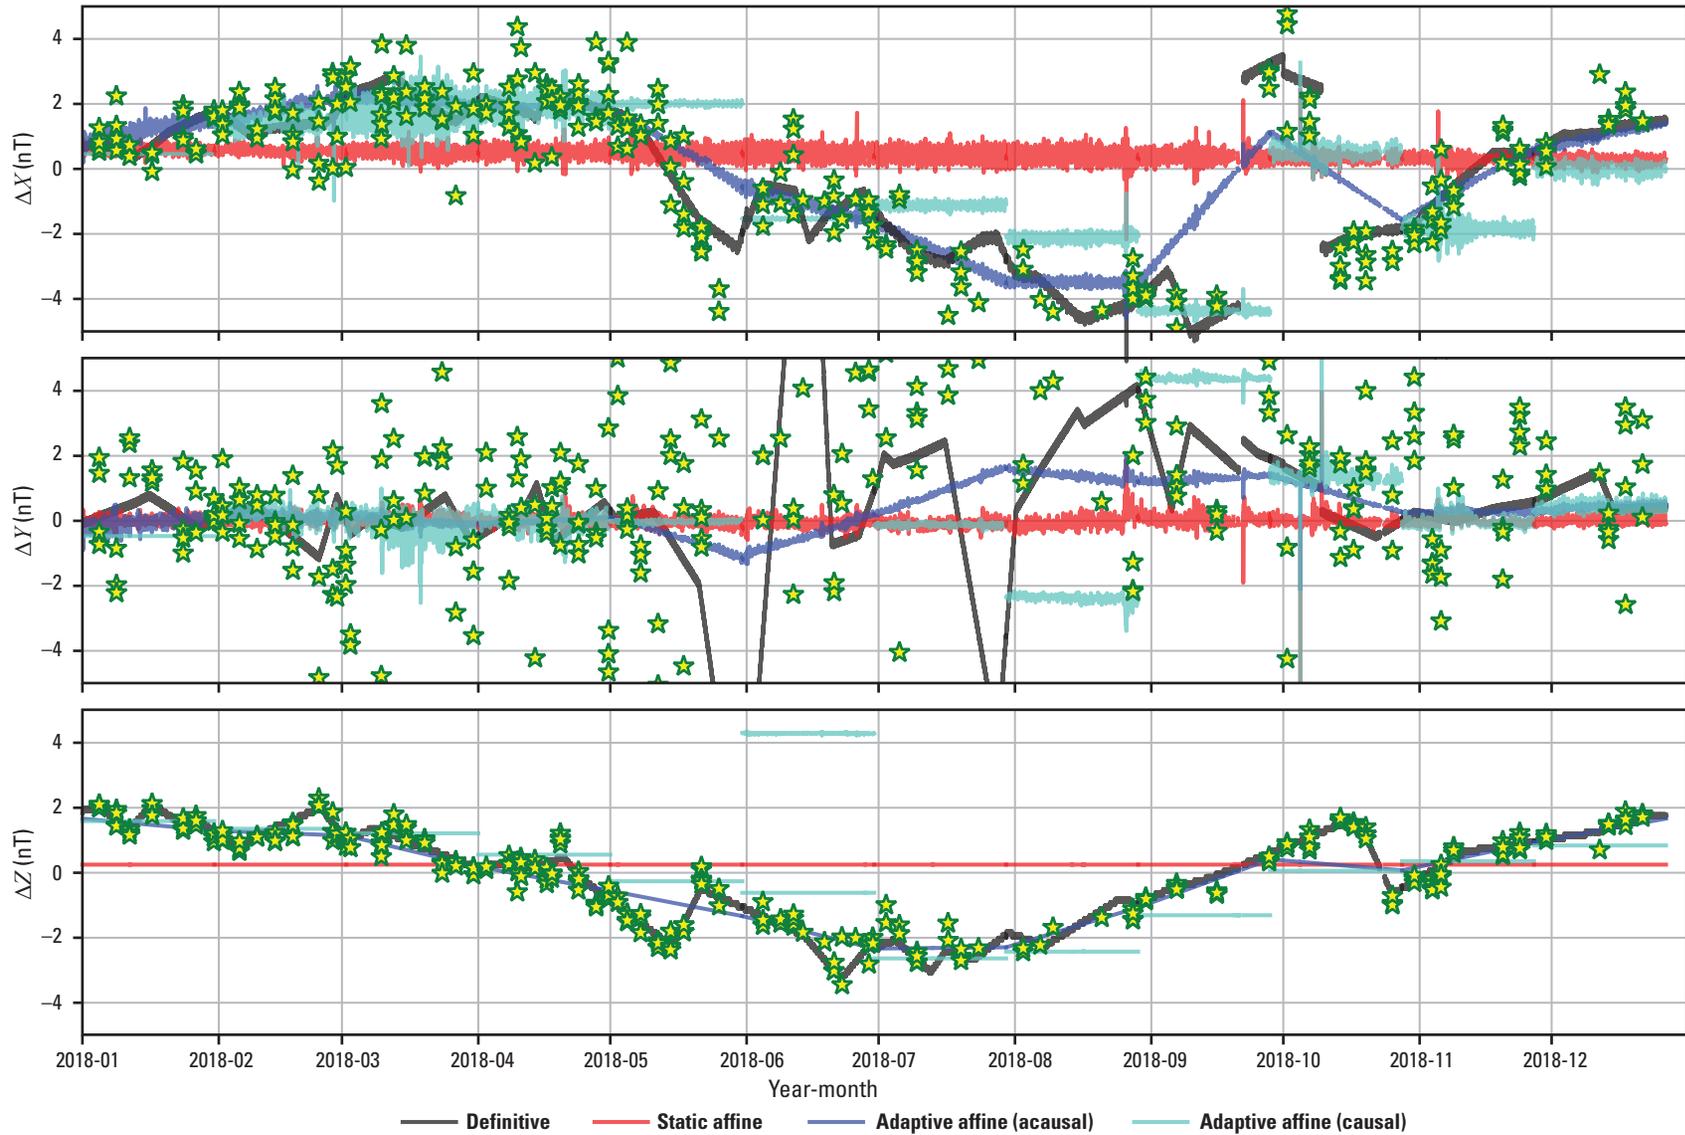


Figure 8. Differences between four raw data correction methods and static traditional baseline adjustments. “Definitive” data are traditional baseline adjustments in which baselines are updated manually to fit each set of absolute measurements. “Static affine” corresponds to an affine transformation in which the matrix is unchanging and constrained to rigid horizontal rotation plus a translation of origins. “Adaptive affine (acausal)” is an adaptive version of the static affine transformation that automatically fits absolute measurements from both before and after each keyframe epoch and uses spherical linear interpolation (Slerp) to interpolate between keyframes. “Adaptive affine (causal)” is a version of the static affine transformation that uses absolute measurements from only before each keyframe epoch. The stars indicate the differences between the static traditional baseline and absolute measurements. (ΔX , difference between variant of corrected X and static baseline-estimated X ; ΔY , difference between variant of corrected Y and static baseline-estimated Y ; ΔZ , difference between variant of corrected Z and static baseline-estimated Z ; nT, nanoTesla)

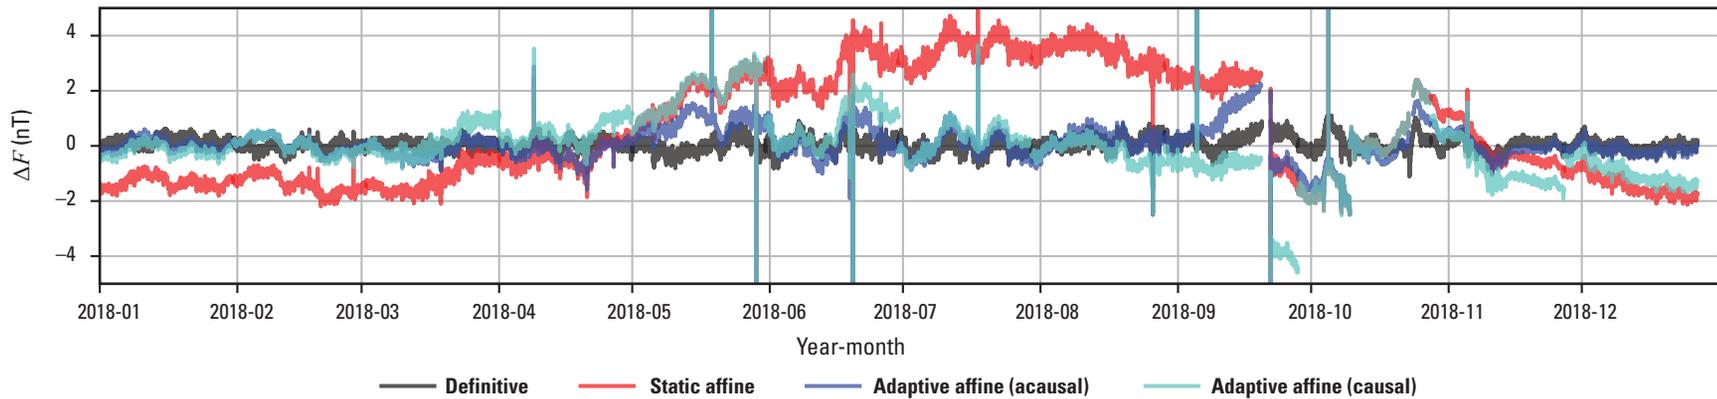


Figure 9. Difference between measured and calculated total field. “Definitive” data are traditional baseline adjustments in which baselines are updated manually to fit each set of absolute measurements. “Static affine” corresponds to an affine transformation in which the matrix is unchanging and constrained to rigid horizontal rotation plus a translation of origins. “Adaptive affine (acausal)” is an adaptive version of the static affine transformation that automatically fits absolute measurements from both before and after each keyframe epoch, and uses spherical linear interpolation (Slerp) to interpolate between keyframes. “Adaptive affine (causal)” is a version of the static affine transformation that uses absolute measurements from only before each keyframe epoch. (ΔF , difference between the XYZ vector amplitude and measured scalar value F , nT, nanoTesla)

techniques. However, these metrics are difficult to reconcile with the data presented in [figure 9](#), especially for the adaptive techniques, which typically appear to be quite close to zero. The reason for this seeming discrepancy is that RMS is notoriously sensitive to outliers because of the “square” in its formulation. The similarly large RMS values for all the affine transformation results are dominated by a handful of spikes in the data. A more robust metric would be mean absolute ΔF deviation (MAD). MAD for Definitive, static affine, adaptive affine (acausal), and adaptive affine (causal) techniques is 0.21, 1.76, 0.44, and 0.76 nT, respectively. The adaptive affine transformation results, especially from the acausal form that considers future absolute measurements, are nearly as good as the Definitive results. The only exceptions are occasional spikes, but upon closer inspection, it becomes clear that these spikes are not just small for Definitive data, but rather the bad data have been removed. In fact, spikes in ΔF are one of the primary criteria used in manual Definitive data processing to identify and remove bad data. Such a spike detector would be simple to code, and may well be incorporated into a production version of the Adjusted data algorithm, but that is a topic for another report.

Summary and Conclusions

The theoretical basis for a new technique to correct raw magnetic variometer data to absolute observations was presented. A brief review of the traditional baseline adjustment technique used by the U.S. Geological Survey (USGS) Geomagnetism Program was contrasted with a newer, more mathematically rigorous, and more physically self-consistent raw data adjustment technique based on affine transformations. A general overview of affine transformations was presented, including techniques to estimate statistically optimal affine matrices from absolute measurements. A more constrained version of the affine transformation matrices was introduced that, roughly, emulates traditional baseline adjustments. This was done partly to limit our departure from the familiar, but also because, given the sparsity and irregularity of absolute observations at most USGS magnetic observatories, limiting the number of free parameters was deemed prudent so that we did not overfit our limited training data.

A synthetic data demonstration was included to verify that the new techniques did indeed perform as intended, but also to familiarize those accustomed to traditional baseline correction methods with the capabilities of affine transformations. Finally, a real data validation was performed using measurements from the USGS BOU magnetic observatory. This validation emulated realistic operational scenarios including (1) a worst case wherein absolute measurements were unavailable to update affine transformations; (2) a real-time scenario wherein no “future” absolute measurements were available; and (3) a Definitive processing scenario wherein “past” and “future” absolute measurements were available to generate monthly keyframes, between which spherical linear interpolation (Slerp) was used to smoothly interpolate to time steps between the keyframes. The first scenario actually met International Real-Time Magnetic Observatory Network (INTERMAGNET) requirements for the BOU magnetic observatory, as well as several observatories that were not presented here. The latter two scenarios met these requirements for all observatories that were known to be functioning nominally, and the final scenario performed nearly as well as Definitive processing, but with no manual intervention.

Two key points that have not been strongly emphasized thus far are (1) the statistical approach to Adjusted data parameter estimation may be a departure from business as usual, but will be critical to future efforts to automate Adjusted data processing because it is more resilient to bad absolute observations that might make it through the usual filters (this is not unique to affine transformations and would be preferred even for traditional baseline estimation); and (2) the affine matrix adjustment method can be used regardless of how the three-axis fluxgate magnetometer is oriented at a given observatory (the same cannot be said for traditional baseline correction methods). As the USGS Geomagnetism Program expands its portfolio of geomagnetic data products and services, accurate, consistent, and timely Adjusted data are an absolute necessity. However, these new products and services require expertise and effort to develop that should not be overburdened with present-day labor-intensive data processing and quality control procedures. The methods presented in this report will go far toward achieving these goals.

Acknowledgments

We thank Adam Ringler, Jeffrey Love, Kristen Lewis, and Jill McCarthy (all of the U.S. Geological Survey) for their insightful comments and constructive feedback to help improve the technical quality and presentation of this report.

References Cited

- Brett, M., 2016, Transforms3d: Git-Hub, Inc. web page, accessed November 2019 at <https://github.com/matthew-brett/transforms3d>.
- Eggert, D.W., Lorusso, A., and Fisher, R.B., 1997, Estimating 3-D rigid body transformations—A comparison of four major algorithms: *Machine Vision and Applications*, v. 9, p. 272–290. [Also available at <https://doi.org/10.1007/s001380050048>.]
- Hitchman, A.P., Crosthwaite, P.G., Jones, W.V., Lewis, A.M., and Wang, L., 2013, Australian Geomagnetism Report 2011: Canberra, Australia, Geoscience Australia, v. 59, Record 2012/72, 229 p.
- Jankowsky, J., and Sucksdorff, C., 1996, Guide for magnetic measurements and observatory practice: Boulder, Colo., International Association of Geomagnetism and Aeronomy, 235 p.
- Love, J.J., Applegate, D., and Townshend, J.B., 2008, Monitoring the Earth’s dynamic magnetic field: U.S. Geological Survey Fact Sheet 2007–3092, 2 p. [Also available at <https://doi.org/10.3133/fs20073092>.]
- Love, J.J., and Finn, C.A., 2011, The USGS Geomagnetism Program and its role in space weather monitoring: *Space Weather*, v 9, S07001, 5 p. [Also available at <https://doi.org/10.1029/2011SW000684>.]
- Millman, K.J., and Aivazis, M., 2011, Python for scientists and engineers: Computing in Science & Engineering, v. 13, no. 2, p. 9–12. [Also available at <https://doi.org/10.1109/MCSE.2011.36>.]
- Oliphant, T.E., 2007, Python for scientific computing: Computing in Science & Engineering, v. 9, no. 3, p. 10–20. [Also available at <https://doi.org/10.1109/MCSE.2007.58>.]
- Shoemake, K., 1985, Animating rotation with quaternion curves, in SIGGRAPH ’85—Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques, San Francisco, Calif., 1985: New York, Association for Computing Machinery, p. 245–254. [Also available at <https://doi.org/10.1145/325334>.]
- Shoemake, K., and Duff, T., 1992, Matrix animation and polar decomposition, in Proceedings of the Conference on Graphics Interface, 18th, Vancouver, B.C., 1992, Proceedings: San Francisco, Calif., Morgan Kaufman Publishers Inc., p. 258–264.
- St-Louis, B., ed., 2012, 2012, INTERMAGNET Technical Reference Manual, version 4.6: Edinburgh, U.K., INTERMAGNET, 92 p. [Also available at <https://www.intermagnet.org>.]
- Thomas, S.W., 1991, Decomposing a matrix into simple transformations, in Arvo, J., ed., Graphics Gems II: Amsterdam, Elsevier Inc., Academic Press, p. 320–323. [Also available at <https://doi.org/10.1016/B978-0-08-050754-5.50069-4>.]

