

ABSTRACT

Title of Dissertation: SAFER GROUNDS: A STUDY OF
LANDMINE DETECTION USING UAV- AND
GROUND-BASED MULTI-MODAL GEOPHYSICS

Heidi Patricia Myers
Doctor of Philosophy, 2024

Dissertation Directed by: Professor Vedran Lekić
Department of Geology
Professor Daniel Lathrop
Department of Geology, Physics

This dissertation addresses the urgent global crisis of landmines, unexploded ordnance (UXO), and explosive remnants of war (ERW) through the lens of multimodal geophysics. Chapter 1 sets the stage by highlighting the humanitarian imperative while underscoring the broader applicability of the developed methods and instruments for shallow critical zone exploration. Unlike conventional engineering-centric approaches, our geoscience-centered methodology offers promising avenues for effectively detecting and characterizing buried hazards. Chapter 2 meticulously examines various geophysical sensors, identifying limitations and proposing innovative solutions. Notably, TetraMag, a novel triaxial magnetic gradiometer, overcomes the deficiencies of single-sensor systems, demonstrating superior sensitivity to small-scale variations in the magnetic field. Chapter 3 delves into the intricate symmetries and invariants of the finite-difference magnetic gradient tensor (FDMGT), elucidating its pivotal role in precise target localization and

parameter estimation within the shallow critical zone. The methodology outlined streamlines data processing and interpretation, laying a robust foundation for UAV-based detection systems. Chapter 4 introduces machine learning techniques, particularly convolutional neural networks (CNNs), as powerful target detection and parameter estimation tools. By synergizing multiple geophysical modalities, these methods enhance our ability to discern subtle anomalies with high accuracy. Chapter 5 proposes a method to mitigate magnetic self-noise in UAV-mounted gradiometers, enhancing data fidelity and spatial coherence. This approach, applicable to various vehicle platforms, further extends the reach of our detection capabilities. In Chapter 6, we integrate and apply these methodologies to a real-world minefield scenario, achieving success in detecting and localizing buried targets. While acknowledging limitations such as payload constraints and computational demands, our findings underscore the versatility and robustness of the developed techniques. This dissertation addresses the pressing humanitarian challenge of landmine detection and advances the broader field of shallow critical zone geophysics. The methodologies and technologies presented here hold promise for diverse applications beyond military contexts, ranging from infrastructure mapping to hydrogeological studies.

SAFER GROUNDS: A STUDY OF LANDMINE DETECTION
USING UAV- AND GROUND-BASED MULTI-MODAL GEOPHYSICS

by

Heidi Patricia Myers

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2024

Advisory Committee:

Professor Vedran Lekić, Chair/Advisor
Professor Daniel Lathrop, Co-Advisor
Associate Professor Karen Prestegard
Associate Professor Mong-Han Huang
Associate Professor Nicholas Schmerr
Professor Martin Rabenhorst

© Copyright by
Heidi Patricia Myers
2024

Dedication

To Eliot, my rock and my heart's truest home, whose love sustains me through every challenge and triumph. To Magneto and Electro, our bundles of endless energy and boundless love, whose presence fills our home with joy. Together, you are the heartbeat of my life, the reason for every smile, and the source of my deepest happiness. Here's to the countless memories we've created and the endless adventures that lie ahead. With all my love and gratitude, ball roll.

Acknowledgments

In embarking on this academic journey, I find myself indebted to a multitude of individuals and institutions whose support and guidance have paved the way for this endeavor.

First and foremost, I extend my deepest appreciation to my esteemed advisors, Ved Lekić and Dan Lathrop. Their steadfast guidance, endless patience, and persistent belief in my abilities have been instrumental in shaping the trajectory of my research. Beyond their scholarly expertise, it is their infectious humor, enduring kindness, and indomitable optimism that have served as beacons of inspiration throughout this journey. To Ved and Dan, you are not only trusted mentors but cherished friends, and for that, I am profoundly grateful.

I extend heartfelt gratitude to my esteemed committee members, Karen Prestegaard, Mong-Han Huang, and Nicholas Schmerr, for their invaluable feedback and encouragement. Each member has enriched this project with their insights and expertise, for which I am sincerely thankful.

This scholarly pursuit would have remained a mere aspiration were it not for the generous support of the National Science Foundation Partnerships for Innovation Research Partnerships (NSF PFI-RP) program and the Department of Defense National Defense Science and Engineering Graduate (DoD NDSEG) Fellowship. I am deeply grateful for the opportunities afforded by these funding sources, allowing me to explore, innovate, and push the boundaries of knowledge in my field. Special thanks to my NDSEG mentor, Joe Conroy, and the Army Research Lab for their invaluable mentorship and support.

I am indebted to Meyer Taffel and Alex Pick-Aluas for their indispensable contributions to the design and fabrication of TetraMag, as well as their enthusiasm and dedication throughout the project. Their expertise and commitment have been pivotal to the success of this endeavor.

To Nolan Ballew, Alessandro Ristelli, and the dedicated staff at IREAP, I extend sincere gratitude for their invaluable assistance and support. Their contributions have been instrumental in bringing this project to fruition.

Heartfelt thanks are due to the UMD CMREC Paintbranch Turfgrass Research Facility and the UMD CMREC Beltsville Facility for their gracious support and provision of essential resources. Special mention is owed to Dave Funk and Kevin Conover for their generosity and assistance in facilitating this research, allowing us to construct a field area on their property, and conducting drone experiments.

I am profoundly grateful to the Geology Department for fostering an enriching academic community and providing sustained support throughout my graduate studies.

To Jasper Baur, Gabriel Steinburg, and the team at the University of Oklahoma's CENFEX facility, I extend sincere appreciation for their hospitality and dedication to our collaborative efforts in humanitarian demining.

My deepest gratitude is reserved for my husband, Eliot, whose unwavering support, boundless love, and unhinged humor have been my constant pillars of strength. To Eliot and our beloved companions, Magneto and Electro, you are my source of endless joy and inspiration; I am eternally grateful for that.

I am indebted to my sister, Ursi, and my parents for their faithful support and invaluable guidance. Their belief in my abilities has been a source of immense encouragement throughout this journey.

To my cherished friends, Heather, Ben, Kiran, Aisha, and Christiana, your uncompromising support and camaraderie have been a source of true strength and joy. I am wholly grateful for your love and friendship.

Finally, I extend thanks to the unsung heroes of this journey: gorilla tape for steadfastly holding things together, Ghost for providing the soundtrack to countless hours of research, and the Creator, from whom all blessings flow.

In closing, I extend wholehearted gratitude to all who have played a role, however large or small, in shaping this academic endeavor. Your contributions have not gone unnoticed; thank you.

Table of Contents

Dedication	ii
Acknowledgements	iii
Table of Contents	vi
List of Tables	x
List of Figures	xi
List of Abbreviations	xxi
Chapter 1: Introduction	1
1.1 Global Landmine and UXO Crisis	1
1.2 The Costs of Inconsistent Detection	3
1.3 Common Geophysical Detection Methods and Limitations	5
1.3.1 Geophysical Detection Methods	5
1.3.2 Single-Sensor Systems	6
1.3.3 Multi-Sensor Systems	20
1.4 Outline of Thesis	21
Chapter 2: Equipment, Sensors, and General Methodology	25
2.1 Geophysical Sensors	25
2.1.1 GPR Technical Specifications	25
2.1.2 Dual IR and Visible Light Photogrammetry	32
2.1.3 LiDAR	34
2.1.4 TetraMag	35
2.2 UAS Integration	40
2.3 General Approach	40
Chapter 3: Enabling Small Anomaly Detection using Finite-Difference Magnetic Gradiometry	45
3.1 Introduction	45
3.1.1 Magnetometry in Geoscience	45
3.1.2 Challenges of Scalar, Vector, and Gradient Magnetometry	47
3.2 Asymmetry of the Finite-Difference Magnetic Gradient Tensor	61
3.3 Invariants of the Magnetic Gradient Tensor	62

3.4	Data	64
3.4.1	Magnetic Gradiometer System (TetraMag)	64
3.4.2	Synthetic Data	65
3.4.3	Field Data	66
3.5	Results	68
3.6	Discussion	74
3.7	Conclusions	77
Chapter 4:	Machine Learning Methods for Detecting and Characterizing Landmines and UXO	80
4.1	Classification and Regression Networks	80
4.2	CNN for Detecting UXO in Radargrams	84
4.3	CNN for Magnetic Source Parameter Estimation	94
4.4	Conclusions	101
Chapter 5:	Method of Self-Noise Reduction for Drone-Based Magnetometry	102
5.1	Introduction	102
5.1.1	Sources of Magnetic Self-Noise	103
5.2	Magnetic Noise Reduction Model	104
5.2.1	Expected $\vec{B}$ Value	106
5.2.2	Motor Current Draw, Speed And Their Time Derivatives	107
5.3	Field Data	109
5.3.1	Inert Minefield	109
5.3.2	Data Preparation	109
5.4	Results	112
5.4.1	Candidate Model A Coefficients	112
5.4.2	Candidate Model B Coefficients	113
5.4.3	Candidate Model C Coefficients	115
5.4.4	Model Performance Comparison	117
5.5	Discussion	123
5.5.1	The Impact of Drone Magnetic Noise on Data Fidelity	123
5.5.2	Sampling Frequency Limitations	124
5.6	Conclusions	126
Chapter 6:	Landmine and ERW Detection using UAV-Based Multi-Modal Geophysics	128
6.1	Field Area Survey and Data Preparation	128
6.2	TetraMag Data	129
6.3	FLIR Data	135
6.4	LiDAR Data	138
6.5	Multi-Modal Anomaly Maps	138
6.6	Discussion	144
6.6.1	Detection Capabilities and Limitations for Individual Sensors	147
6.6.2	Detection Capabilities and Limitations of Multi-Modal Maps	150
6.7	Conclusions	151

Chapter 7: Conclusions and Future Research	153
7.1 Chapter Conclusions	153
7.2 Lessons Learned	158
7.2.1 Sourcing and Operating High-Payload Drones	158
7.2.2 The Importance of Appropriate Sampling Frequency	158
7.2.3 The Importance of Magnetometer Calibration	159
7.2.4 Limiting Factors for Drone-Based Geophysics	159
7.2.5 The Challenge of Relying on Open-Source Tools	160
7.3 Future Research	160
7.3.1 UAV-Optimized Triaxial Fluxgate Magnetometers	160
7.3.2 Utilizing Flight Controller for Instrument Control	161
7.3.3 Quantifying Optimal Survey Conditions and Correcting FLIR Data with an Albedometer	161
7.3.4 Integrating Mid-Wavelength FLIR Camera for Gas Leak Detection	162
7.3.5 Drone-Based Software-Defined GPR	162
7.3.6 Shallow Critical Zone Moisture Mapping with Drone-Based SDR GPR and FLIR	163
7.3.7 Detection and Mapping of Environmental Contaminants	163
7.3.8 Mapping of Oxidation Fronts for Magnetic Minerals	164
7.3.9 Flight Path Optimization using FDMGT Invariants	164
Appendix A: Exploring Principal Component Analysis, Convolutional Neural Networks, and Scattering Transforms in a Rock Image Classifier	166
A.1 Machine Learning in Geosciences	167
A.1.1 Multi-Layer Perceptron (MLP)	168
A.1.2 Principal Component Analysis (PCA)	168
A.1.3 Convolutional Neural Networks (CNNs)	169
A.1.4 Scattering Transform (ST)	170
A.2 Data Organization and Methods	171
A.2.1 PCA Methods	171
A.2.2 CNN Methods	172
A.2.3 ST Methods	172
A.2.4 t-Distributed Stochastic Neighbor Embedding (t-SNE) for Visualization	174
A.2.5 MLP Methods	174
A.3 Results	176
A.3.1 PCA Results	176
A.3.2 CNN Results	179
A.3.3 ST Results	183
A.4 Discussion	195
A.4.1 PCA: Grayscale vs Color	195
A.4.2 CNN: Grayscale vs Color	197
A.4.3 ST: Grayscale vs Color	199
A.5 Conclusions	209
Appendix B: Standard Operating Procedures and Checklists	211

B.1	TetraMag SOP	211
B.2	FLIR Duo Pro R	214
B.3	GeoSLAM Zeb Horizon	215
	Bibliography	220

List of Tables

5.1	Target specifications for Glenn Dale, MD, field site. Abbreviations: D (diameter), H (height), L (length), W (width), AT (anti-tank), AP (anti-personnel), IED (improvised explosive device), EM (electromagnetic), PDM (pursuit deterrent munition).	110
5.2	Average order of magnitude variance reduction by component for each candidate model.	119
5.3	Root mean squared error (RMSE) between different components of $\mathbf{G}$ using both uncorrected and corrected TetraMag data for each model.	121
6.1	Detected ferrous target measured and predicted burial depths. Predicted depths are calculated using the Frobenius norm ratio $\ \mathbf{G}_A\ _F/\ \mathbf{G}_S\ _F$	134

List of Figures

1.1	Landmine, UXO, and ERW Contamination and Casualties in 2022. (Adapted from 1.1)	2
1.2	Types of Landmines and Common Specifications.	3
1.3	Common Landmine Detection Techniques. Geophysics-based techniques are in yellow. (Adapted from [1])	5
1.4	Precession around an external magnetic field. A proton's angular momentum, L , and parallel magnetic field vector, $\vec{m}$, changes in the presence of an external magnetic field, $\vec{B}$, causing a precession around $\vec{B}$ proportional to the proton-gyromagnetic ratio. (Adapted from [2])	8
1.5	Fluxgate magnetometer diagram. Primary coils and oppositely-wound secondary coils surround two ferromagnetic cores. Driving current through the primary coils will induce an opposite magnetic field from the primary coil. When the cores are aligned with an external magnetic field, the resulting asymmetry in the secondary coil will produce a voltage proportional to the magnitude of the external field. (Source: [2])	9
1.6	Ring-Core Fluxgate Magnetometer Diagram. Source: [3]	10
1.7	Common GPR antenna geometries and 3D radiation patterns (right side of each antenna). (a) Circular Plate Bow-Tie. (b) Triangular Plate Bow-Tie. (c) Horn. (d) Vivaldi. (e) Archimedean Spiral. (e) Equiangular Spiral.	13
1.8	A typical GPR survey setup. The lawnmower-style commercial GPR unit sends EM waves into the ground, producing pulse echoes (A-scans). As the GPR moves across a survey line, the accumulated A-scans combine into a B-scan, known as a radargram, showing the subsurface profile in 2-D. Surveying with sufficient tie points allows for the construction of a C-scan, the 3-D volume representation of the survey area.	14
1.9	Example of thermal (IR) imagery used to detect buried plastic and metal mines. (Adapted from Figure 6 of [4])	16
1.10	Example of using LiDAR to map minefield features. (Adapted from Figure 7 of [5])	19
2.1	GSSI SIR 4000 pushcart with 400 MHz antenna.	26
2.2	Condensed target grid at humanitarian minefield in Pawnee, OK. (Figure 3 from [6])	27
2.3	Example processed radargrams from GPR inert minefield study. The top radargram has clear hyperbolas visible (indicated by dashed circles) and good localization. The bottom radargram shows some hyperbolas, but localization is poor due to the GPR being lifted multiple times during the run (indicated by black arrows).	29
2.4	Expected target and GPR hyperbola locations.	30

2.5	Summary of target types, their dimensions, and burial depths used for height-varying GPR experiment.	31
2.6	Radargrams from height-varying GPR survey over landmines and analog UXO. .	33
2.7	FLIR Duo Pro R dual visible light and thermal camera.	34
2.8	GeoSLAM Zeb Horizon 3-D LiDAR.	35
2.9	TetraMag initial design (a) and with added structural supports (b).	37
2.10	TetraMag field test setup using a height-adjustable nonmagnetic aluminum scaffold.	38
2.11	TetraMag v2, designed as landing gear for the Freefly Alta X UAV.	39
2.12	Alta X flying time vs payload weight. Source: [7].	41
2.13	Freefly Alta X with integrated drone design (top) and implementation (bottom). .	42
2.14	General approach summary. Geophysical sensors (FLIR Duo Pro R, GeoSLAM Zeb Horizon, TetraMag) are coupled with telemetry sensors (measuring motor RPM, motor current draw, temperature, humidity) on a quadcopter UAV (Freefly Alta X).	43
2.15	Sensor Modalities, Controls, and Communications.	44
3.1	Coherence between magnetic reference stations during geomagnetically quiet conditions, Kp index 0 (a,c), and a G3 geomagnetic storm, Kp index 7 (b,d), at locations across the United States: Boulder, CO (a), Tucson, AZ (b), Fredericksburg, VZ (c), and NASA Stennis Space Center, MS (d). In quiet conditions, all components' coherence is high at low frequencies and falls off rapidly at frequencies $> 50\text{ mHz}$. Vertical and total field coherence falls off during geomagnetic storms, such as in quiet conditions, but North and East components show enhanced coherence at higher frequencies.	49
3.2	Expected magnetic total field anomaly (B_t) of a dipole-like anti-tank (AT) mine buried at 15 cm depth measured from various sensor heights, compared to typical amplitudes of geomagnetic storms and nearby DC Metro systems (shaded areas) and the noise floor of magnetometer used in this study, MEDA FVM400 (dashed line). At heights $> 0.5\text{ m}$, geomagnetic storms and DC Metro signals can obfuscate the AT mine, producing false flags.	50
3.3	Signal error as the percentage of expected magnetic anomaly for a metal AT mine buried at 15 cm depth, measured with a triaxial fluxgate magnetometer from 0.5 m height, as a function of the pointing (heading) error of the sensor. $B_x = 150\text{ nT}$, $B_y = 200\text{ nT}$, $B_z = 400\text{ nT}$, declination = -10.5° , inclination = 67.12° (matching College Park, MD). Very small pointing errors produce spurious magnetic signals large enough to obscure those from an AT mine, increasing the chances of false flags.	53
3.4	(a) Schematic of a regular tetrahedron-based triaxial magnetic gradiometer. 1, 2, 3, and 4 are the positions of triaxial fluxgate magnetometers, each at the vertices of a regular tetrahedron. The side length is labeled as d . Each magnetometer will have x, y, and z components (B_{1x} , B_{1y} , B_{1z} , B_{2x} , etc). (b) TetraMag magnetic gradiometer system. MEDA FVM400 triaxial fluxgate magnetometer sensor heads are in 3D printed enclosures at the four vertices of a regular tetrahedron with $d = 0.5\text{ m}$ side length.	56

3.5	The $\hat{x}_3$ finite difference of the $\hat{x}_1$ component of the magnetic field ($\mathbf{G}_{13}$, black solid line) and the $\hat{x}_1$ finite difference of the $\hat{x}_3$ component of the magnetic field ($\mathbf{G}_{31}$, red solid line) are poorly approximated when only the first term in equations 3.10 (black dashed), 3.16 (red dashed) is used. However, the inclusion of second-order terms in ϵ (dot-dashed lines) improves accuracy. Note that the dashed lines are parallel because of the symmetry of $\vec{\nabla} \vec{B}$	60
3.6	Symmetric (red) and antisymmetric (blue) components of FDMGT exhibit different behavior with increasing gradiometer side lengths. $\ \mathbf{G}_{\Gamma 30785}\ _F$ grows with the square of the side length (dotted line), while $\ \mathbf{G}_{\Gamma 30803}\ _F$ remains unchanged so that their ratio also grows with the square of the side length. For small side lengths, the antisymmetric component is negligible but approaches the symmetric component for realistic gradiometer sizes.	62
3.7	Survey geometry and example schematic. (a) Magnetic field lines for a vertical dipole buried at 15 cm depth and TetraMag, to scale (lower left), with orientations of B_x , B_y , and B_z for magnetometer 1. (b) Lawnmower-style survey patterns followed by TetraMag, with gradients, $\mathbf{G}_{\Gamma 30790}$, referenced to its centroid (black star). (c) An example visualization for one of the $\mathbf{G}_{\Gamma 30790}$ invariants, with color of each pixel representing a chosen $\mathbf{G}_{\Gamma 30790}$ invariant at that location.	67
3.8	Maps of the eigenvalues (columns 1-3) and invariants (columns 4-11) calculated from synthetic data at two TetraMag orientations for the full IMGT and FDMGT, $\mathbf{G}_{\Gamma 30822}$, as well as the latter's symmetric, $\mathbf{G}_{\Gamma 30803}$, and antisymmetric, $\mathbf{G}_{\Gamma 30785}$, parts. Eigenvalues λ_1 , λ_2 , and λ_3 are shown in columns 1-3. Determinant variations, $ \mathbf{G}_{\Gamma 30822} $ and $ \mathbf{G}_{\Gamma 30803} $, are shown in columns 4-5. Frobenius norm variations, $\ \mathbf{G}_{\Gamma 30822}\ _F$, $\ \mathbf{G}_{\Gamma 30803}\ _F$, and $\ \mathbf{G}_{\Gamma 30785}\ _F$, are shown in columns 6-8. Variations of eigenvalue combinations 1 (λ_{C1}) and 2 (λ_{C2}) and NSS are shown in columns 9-11.	70
3.10	Maps of eigenvalues (columns 1-3) and invariants (columns 4-11) for field FDMGT data (Field) collected by TetraMag and synthetic FDMGT data (Synth.) calculated using TetraMag's geometry, tilt, and local magnetic conditions. Eigenvalues λ_1 , λ_2 , and λ_3 are shown in columns 1-3. Determinant variations, $ \mathbf{G} _f$ and $ \mathbf{G} _S$, are shown in columns 4-5. Frobenius norm variations, $\ \mathbf{G}_{\Gamma 30822}\ _F$, $\ \mathbf{G}_{\Gamma 30803}\ _F$, and $\ \mathbf{G}_{\Gamma 30785}\ _F$, are shown in columns 6-8. Eigenvalue combinations 1 (λ_{C1}) and 2 (λ_{C2}) and NSS variations are shown in columns 9-11.	72
3.9	Average B_x , B_y , and B_z values for synthetic data and corresponding field data measured by TetraMag at 55 cm height. The dipole target is indicated by a black asterisk and buried 15 cm in a test sand bed in College Park, MD. Synthetic data match the field magnetic inclination and declination values. Noise from magnetic clutter on the test scaffold has been piecewise removed from the field B_y values, ghosting from which can be seen in the plot.	73

3.11	Ratio of $\ G_{\Gamma 30785}\ _F / \ G_{\Gamma 30803}\ _F$ at two flying heights, 0.5 <i>m</i> (top row), corresponding to a ground-based survey, and 1.5 <i>m</i> (bottom row), corresponding to a low-altitude drone-based survey. The source is a magnetic dipole that is dipping (left), horizontal (center), or vertical (right). While a vertical magnetization produces the most symmetric $\ G_{\Gamma 30785}\ _F / \ G_{\Gamma 30803}\ _F$ pattern around the target location, dipping orientations produce more complicated, lobate patterns, but still peak above the target.	76
4.1	Architecture of a feed-forward multi-layer perceptron (MLP) network. (Source: [8])	81
4.2	Typical nonlinear activation functions. (a) Sigmoid, (b) Tanh, (c) ReLU, and (d) LeakyReLU. (Source: [9])	83
4.3	CNN architecture example. An image undergoes two convolution and max pooling functions before the results are flattened, passed through a fully connected MLP, then classified into 2 categories at the output layer. (Source: [10])	84
4.4	Analog UXO Field Site. Pictures on the right show the analog UXO devices made of metal (top), plastic (middle), and ceramic (bottom).	85
4.5	Example pre-processed GPR radargrams from the UXO experiment. The horizontal axis of the radargram represents 5 m length, and the vertical axis represents 0-2 m depth. The upper panel contains UXO (location indicated with a red box and zoomed inset). The lower panel contains no UXO.	86
4.6	CNN Architecture for UXO Detection in Radargrams. Input images with 90 pixel width and 40 pixel height are fed through a series of convolution and pooling layers, followed by 2 fully connected layers before the output layer. The first convolution layer has 32 filters and a kernel size of 5 x 5. The second and third convolution layers have 48 filters, with 4 x 4 and 3 x 3 kernels, respectively. The first fully connected layer has 64 neurons, the second fully connected layer has 16 neurons, and the output layer has 1 neuron.	88
4.7	Model weight histograms for each CNN layer. The distribution of true model weights is displayed at each epoch step, where epoch 0 is the first epoch, and epoch 3 is the fourth epoch. Changes in shape or amplitude between epochs suggest the model is using that layer to learn successfully. Changes in shape and amplitude are present for all layers, supporting the use of all in the final CNN design.	90
4.8	Examples of input radargram image changes after each epoch. (a) shows an example image containing UXO, while (b) does not contain UXO. Important features are kept by the CNN after each epoch, while unimportant features are discarded.	91

4.9	Loss and accuracy for training and testing phases of the UXO CNN for 10 consecutive runs. Panels (a) and (b) show loss values for training and testing, respectively. While the rates vary, each run shows decreasing loss with increasing epoch. Panels (c) and (d) show accuracy values for training and testing, respectively. Both training and testing accuracies increase with increasing epoch, but the testing accuracy lags slightly behind training accuracy. This is likely a function of the model's weights and biases changing drastically in epochs 1 and 2. Maximum training accuracy values range between 86-99%, while testing accuracy values range between 84-95%.	92
4.10	UXO CNN Confusion Matrix and Classification Report. The upper panel shows the confusion matrix for the CNN, where 0 represents "No UXO Detected" and 1 represents "UXO Detected". A false negative is the most dangerous case for a de-miner; the CNN was able to produce less than 5% false negatives for the 1000 radargrams used for validation. The F-1 score represents the weighted average of precision and accuracy for the model, where 1 is a perfect model.	93
4.11	CNN architecture for estimation source parameters using FDMGT invariant maps.	95
4.12	Magnetic moment estimation from source parameter CNN. The red line is a 1:1 reference line.	97
4.13	Magnetic data of interest as included CNN layers vs. magnetic moment estimation accuracy.	98
4.14	Burial depth estimation from CNN. The red line is a 1:1 reference line.	99
4.15	Magnetic data of interest as included CNN layers vs. burial depth estimation accuracy.	100
5.1	Sources of magnetic noise on a drone. (a): current draw from DC power bus to motors. (b): Earth's ambient magnetic field. (c): vibration of the platform. (d) motion of permanent magnets inside BLDC motors. (e): payload instrumentation.	104
5.2	Magnetic noise reduction workflow.	105
5.3	Inert Minefield and target locations near Glenn Dale, MD. The photo inset shows the Alta X drone with our geophysical sensor payload.	111
5.4	Comparison between no filter and 0.25 Hz low-pass filtered data. The filtered data has better spatial coherence as it suppresses vibration from the drone flying.	112
5.5	Model A coefficients. Input data and model variables have been normalized to std = 1. The four magnetometers of TetraMag are labeled "Mag1", "Mag2", "Mag3", "Mag4".	114
5.6	Model B coefficients. Input data and model variables have been normalized to std = 1. The four magnetometers of TetraMag are labeled "Mag1", "Mag2", "Mag3", "Mag4".	116
5.7	Model C coefficients. Input data and model variables have been normalized to std = 1. The four magnetometers of TetraMag are labeled "Mag1", "Mag2", "Mag3", "Mag4".	118

5.8	Model performance summary. (i) shows the root mean squared error for each component by model. "Exp B" uses only the geometrically-compensated values. (ii) shows an example signal from Magnetometer 4, z-component (${}^C B_{4z}$), with the values predicted by Model C. (iii) shows the anomalies from Model C. These values match the expected anomaly ranges for landmines, UXO, and ERW. (iv) shows predicted B_{4z} vs measured B_{4z} . The red line is a 1:1 reference.	120
5.9	Average B_x , B_y , and B_z anomalies by magnetic noise model, with target locations indicated. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.	122
5.10	The determinant of the symmetric part of the full magnetic gradient tensor ($ G_S $) for all models. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.	125
6.1	Eigenvalue anomaly maps from a drone-based survey of the inert minefield. A white "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A white "o" symbol indicates GNSS locations of non-ferrous targets and control holes.	131
6.2	Determinant and Frobenius norm anomaly maps from a drone-based survey of the inert minefield. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.	132
6.3	Frobenius norm ratio $\ G_A\ _F/\ G_S\ _F$ map of the inert minefield (top) and the corresponding predicted vs measured depth values (bottom). The red line in the bottom panel is a 1:1 reference line.	133
6.4	Eigenvalue combinations and normalized source strength anomaly maps from a drone-based survey of the inert minefield. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.	136
6.5	Visible light orthomosaic of the larger field area. Survey markers appear as white boxes in the photo. Examples of detected anomalous areas are shown in the inset boxes. The top left inset shows target B7. The bottom center and right insets show areas where oval-shaped patches of soil differ from the surrounding vegetation (targets C3 and C1, respectively).	137
6.6	Visible light orthomosaic of the inert minefield. Anomalous features are in black circles.	138
6.7	Thermal orthomosaic of the field area. Temperature data are relative, where 0 is the coldest, and 1 is the hottest spot. White squares are ground control points. . .	139
6.8	Thermal orthomosaic of the inert minefield with anomalous features circled in white.	140
6.9	Thermal orthomosaic of the inert minefield with anomalous features isolated and circled in white.	140

6.10	Processed GeoSLAM Zeb Horizon LiDAR data. (a) shows an oblique view, and (b) a nadir view of the survey area. The coordinate system is referenced to the drone takeoff location, not the survey markers.	141
6.11	Georeferenced LiDAR map with anomalies circled in white.	142
6.12	Visible light orthomosaic draped on the LiDAR return intensity map. Anomalies present only in LiDAR are circled in white. Anomalies in both LiDAR and visible light are circled in pink.	143
6.13	Thermal and visible light orthomosaics draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow, and anomalies present in LiDAR, visible, and thermal orthomosaics are circled in pink.	143
6.14	Visible light, thermal, and λ_2 invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.	144
6.15	Visible light, thermal, and $ G_S $ invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.	145
6.16	Visible light, thermal, and $\ G_S\ _F$ invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.	145
6.17	Visible light, thermal, and λ_{C1} invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.	146
6.18	Visible light, thermal, and λ_{C2} invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.	146
6.19	Visible light, thermal, and NSS invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.	147
6.20	Visible light, thermal, λ_{C2} , $\ G_S\ _F$, $ G_S $, and NSS invariant maps draped on the LiDAR return intensity map.	148
A.1	Convolutional Neural Network (CNN) architecture. Each of the 23998 rock images is fed in and traverses through two convolution layers, with max pooling functions between them, then flattened at the output layer. The flattened outputs are then fed into the MLP for classification.	172

A.2	Multilayer Perceptron (MLP) architecture. Outputs from either PCA, CNN, or ST are used as input to the MLP, where they traverse through two fully-connected layers containing 33 and 48 nodes, respectively, before reaching the output layer, containing 6 nodes corresponding to the 6 rock types (andesite, gneiss, marble, quartzite, rhyolite, and schist). Positive edges between nodes are shown in yellow, and negative edges are shown in orange	175
A.3	Accuracy and Loss of PCA-based Rock Classifier (Grayscale Images) (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	177
A.4	Confusion matrix of PCA-based Rock Classifier (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	178
A.5	Accuracy and loss for PCA-based Rock Classifier (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	179
A.6	Confusion matrix for PCA-based Rock Classifier (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	180
A.7	Accuracy and loss for CNN-based Rock Classifier (Grayscale Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	181
A.8	Confusion matrix for CNN-based Rock Classifier (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	182
A.9	Accuracy and loss for CNN-based Rock Classifier (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	183
A.10	Confusion matrix for CNN-based Rock Classifier (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	184
A.11	Accuracy and loss for ST-based Rock Classifier, $L = 2$ (Grayscale Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	186
A.12	Confusion matrix for ST-based Rock Classifier, $L = 2$ (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	187

A.13 Accuracy and loss for ST-based Rock Classifier, $L = 8$ (Grayscale Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	188
A.14 Confusion matrix for ST-based Rock Classifier, $L = 8$ (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	189
A.15 Accuracy and loss for ST-based Rock Classifier, $L = 2$ (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	191
A.16 Confusion matrix for ST-based Rock Classifier, $L = 2$ (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	192
A.17 Accuracy and loss for ST-based Rock Classifier, $L = 8$ (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.	193
A.18 Confusion matrix for ST-based Rock Classifier, $L = 8$ (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.	194
A.19 t-SNE projection of PCA scores (Grayscale Images).	196
A.20 t-SNE projections of PCA scores for 3-channel input images. PCA was performed on red, green, and blue channels separately, then combined to create the 3-channel principal components. t-SNE representations for the separate channels are shown in the upper portion, and the representation for the combined principal components is shown below	198
A.21 t-SNE projections for CNN extracted features after the second convolution layer (a) and before the last layer (b) for grayscale input images.	200
A.22 t-SNE projections for CNN extracted features after the second convolution layer (a) and before the last layer (b) for 3-channel color input images.	201
A.23 t-SNE plot for summary ST coefficients with $J = 4$, $L = 2$, $M = 2$ (Grayscale Images).	203
A.24 t-SNE plot of summary ST coefficients with $J = 4$, $L = 8$, $M = 2$ (Grayscale Images).	204
A.25 t-SNE projections of ST summary coefficients for 3-channel input images, $J = 4$, $L = 2$, $M = 2$. STs and summary coefficients were calculated on red, green, and blue channels separately and then combined to create the 3-channel ST summary coefficients. t-SNE representations for the separate channels are shown in the upper portion, and the representation for the combined coefficients is shown below.	206
A.26 t-SNE projections of ST summary coefficients for 3-channel input images, $J = 4$, $L = 8$, $M = 2$. STs and summary coefficients were calculated on red, green, and blue channels separately and then combined to create the 3-channel ST summary coefficients. t-SNE representations for the separate channels are shown in the upper portion, and the representation for the combined coefficients is shown below.	208

B.1	FLIR Duo Pro Operation Buttons.	215
B.2	GeoSLAM Zeb Horizon SOP	216

List of Abbreviations

AP	Anti-Personnel Landmine
AT	Anti-Tank Landmine
BART	Bay Area Rapid Transit
COTS	Commercial Off-the-Shelf
DC	Direct Current
EM	Electromagnetic
EMI	Electromagnetic Induction
ERW	Explosive Remnants of War
FDMG	Finite-Difference Magnetic Gradiometry
FDMGT	Finite-Difference Magnetic Gradient Tensor
FMCW	Frequency-Modulated Continuous Wave
GNSS	Global Navigation Satellite System
GPR	Ground Penetrating Radar
IED	Improvised Explosive Device
IR	InfraRed
LiDAR	Light Detection and Ranging
MGT	Magnetic Gradient Tensor
ML	Machine Learning
NMO	Normal Move-Out
NQR	Nuclear Quadrupole Resonance
RPM	Revolutions Per Minute
RTK	Real-Time Kinematic
SFCW	Stepped-Frequency Continuous Wave
UAV	Unmanned Aerial Vehicle
UAS	Unmanned Aerial System
USD	U.S. Dollars
UXO	Unexploded Ordnance

Chapter 1: Introduction

1.1 Global Landmine and UXO Crisis

There are over 110 million landmines, unexploded ordnance (UXO), and explosive remnants of war (ERW) buried worldwide that cause thousands of injuries and/or deaths every year [11] (see Figure 1.1). In addition to the 110 million already emplaced landmines, there are roughly 250 million stockpiled, and 2.5 million new landmines are buried each year, with recent conflicts resulting in a tenfold increase in casualties in places such as Ukraine [12, 13]. Landmines are made from various materials, including metal, plastic, ceramics, wood, and glass, and their placement occurs across diverse landscapes, including arid deserts, humid continental, and damp tropical zones [14]. Improvised explosive devices (IEDs) are also part of the modern warfare scene, largely due to their low cost, fast construction time, and easily concealed nature [13, 15].

There are two main categories of landmines: anti-tank (AT) and anti-personnel (AP) (see Figure 1.2). While landmines can be traced back to the Roman Empire, AT designs were not invented until the 1800s and have been in continual use ever since [16]. AP landmines were first used in World War I, and were primarily made of metal; in response, the military relied heavily on metal detectors for detection and removal efforts [17]. World War II saw an increase in landmines made of other materials, such as plastic and ceramic, to avoid detection by handheld

metal detectors [18].

While peacetime agreements such as the 1997 Mine Ban Treaty lay the framework for ending the use, production, stockpiling, and transfer of AP mines and for destroying stockpiles and clearing mined areas, the reality is that landmines and ERWs continue to plague the globe [13, 19]. The worldwide proliferation of buried explosives creates a long-lasting threat of victim-activated detonation along with social and economic harm to the affected re-

gions. In Lebanon, for example,

the loss of agricultural resources

due to landmine contamination amounts to USD \$126.7 million in direct losses, potentially four times that amount when accounting for indirect costs [20]. Since managing the landmine and UXO threat is now beyond the reach of international sanctions, remediation efforts must be effective and efficient.




Figure 1.1: Landmine, UXO, and ERW Contamination and Casualties in 2022. (Adapted from 1.1)




Figure 1.2: Types of Landmines and Common Specifications.

1.2 The Costs of Inconsistent Detection

Inconsistent ERW and UXO detection techniques have made remediation efforts slow, dangerous, and stunted agricultural development in countries most impacted by minefields. Mine remediation remains incredibly expensive — it costs an average of USD \$3 to make a simple landmine and an average of USD \$1000 to clear it [21]. Locations such as the Falklands have dealt with minefield contamination since 1982, and are still not fully cleared of landmines [22]. Boggy environmental conditions inhibit traditional geophysical surveying, including ground penetrating radar (GPR) and magnetometry [23]. Slow remediation has devastating consequences for the safety of people living near the minefields and prevents citizens from using it for agricultural development, directly affecting the local economy and ability to feed its inhabitants. In countries such as Colombia, where landmine-related casualties are increasing rapidly, minefield contami-

nation has slowed the battle against the drug trade, exacerbating an already devastating narcotics crisis [24].

A cost-efficient, easily deployed, terrain-independent, and automated approach to detection and remediation is needed. One proposed solution is to combine geophysical remote sensing with UAV (Unmanned Aerial Vehicle) technology, but current aerial solutions deploy single-sensor capabilities that do not adjust to environmental conditions. GPR, for example, is best suited for environments where soil moisture content is low and the substrate is homogeneous, which is not common for many minefields. Additionally, the utility of GPR when used in diverse soil and environmental conditions has yet to be fully characterized, especially for UAV use cases [25]. Machines designed to clear minefields by detonating as many mines as possible are harmful to the environment, as detonation leads to pollutant release and soil degradation or deforestation, which may affect entire populations of species by altering food chains and damaging habitats [26]. Commercially available single-sensor demining unmanned aerial systems (UAS) often yield false flags, while current multi-sensor systems are land-based, cumbersome to transport to remote sites, terrain limited, and require hazardous on-site participation and exposure to potential blast radii. Addressing the global UXO crisis using purely engineering techniques (demining machines) or purely scientific techniques (traditional single-sensor and land-based multi-sensor surveying) has not worked. A multi-disciplinary approach that relies heavily on geoscience and our understanding of geophysics may be the key to solving this problem and saving thousands of lives. This work aims to investigate the advantages and limitations of a multi-modal geophysical sensor-fused, non-invasive aerial package and explore the extent to which novel data processing techniques and machine learning may be leveraged to achieve optimal anomaly detection in the shallow critical zone.




Figure 1.3: Common Landmine Detection Techniques. Geophysics-based techniques are in yellow. (Adapted from [1])

1.3 Common Geophysical Detection Methods and Limitations

1.3.1 Geophysical Detection Methods

Figure 1.3 summarizes the most common landmine detection techniques in use today [1]. Geophysical principles and practices that are most applicable are in yellow and include magnetics, GPR, microwave radiometry, electrical impedance tomography, infrared (IR), visible light, light detection and ranging (LiDAR), nuclear quadrupole resonance (NQR), sound and ultrasound, and acoustic/seismic. Of these, the non-invasive techniques that may be integrated as a UAV payload include magnetics, GPR, visible light, LiDAR, and NQR. To understand which methods were best suited for this project, exploring the operational theory, advantages, and limitations of commercially available single-sensor and multi-sensor systems is useful.

1.3.2 Single-Sensor Systems

Magnetometers (proton precession and fluxgate types), GPR, IR, and LiDAR are the most widely used single-sensor systems for landmine and ERW detection.

1.3.2.1 Magnetometers

Terrestrial magnetometers measure the magnetic field strength at a particular location on Earth, which results from a combination of (a) the slow-varying field that originates from magnetohydrodynamic interactions in the outer core, (b) the fast-varying field that originates from ionospheric currents, (c) anomalies that result from geophysical targets, and (d) variations that result from anthropogenic activities [27]. Geophysical targets vary in spatial scale from kilometers (for petroleum and mining materials) to centimeters (for archeological and UXO/landmine materials), and the intensity of the anomaly will depend on the magnetic properties of the material and its relative abundance. Anthropogenic activities that can result in magnetic field variations include the use of DC-powered train systems such as the Bay Area Rapid Transit (BART) electric railway in San Francisco, which produces both electromagnetic noise and Pc3 (7-100 mHz) geomagnetic pulses [28, 29]. Earth's magnetic field has dipolar and nondipolar components, and westward drift of approximately 0.2° in longitude occurs annually [30]. The goal of a magnetometer-based study, then, is to measure the compounded magnetic field (total-field intensity), separate the contributions of the ambient field from anomalies associated with geophysical targets, and model the observed field, using complementary geologic information to select the most likely candidate model for interpretation. Magnetometers exploit the ferrimagnetic properties of geologic and anthropogenic materials. When a ferrimagnetic material (such as a landmine with metal casing)

is placed in an external magnetic field, induced magnetization occurs, dissipating once the external field is removed [31]. The two most common instruments used for field-based magnetic measurements and landmine/ERW surveying are proton precession magnetometers and fluxgate magnetometers [32–38].

Proton Precession Magnetometers Proton precession magnetometers estimate the scalar magnetic field by measuring the Larmor precession frequency of a proton magnetic moment around an ambient magnetic field [2]. The sensor head of commercially available proton precession magnetometers is a cylinder filled with hydrogen-rich fluid (typically decane or water) surrounded by a coil and connected to a power supply, amplifier, and frequency-measuring electronics. When a current is applied to the coil, a magnetic field, $\vec{B}$, that is orders of magnitudes larger than Earth’s ambient field is produced parallel to the coil axis [31]. The protons in the cylinder’s liquid act as tiny spinning dipoles that align themselves to the direction of the coil’s magnetic field. The ratio of the proton’s spin angular momentum, L , and parallel magnetic moment vector, m_p is the well-known proton-gyromagnetic ratio, γ_p [39].

$$\gamma_p = \frac{m_p}{L} = 2.6752218744 \times 10^8 s^{-1}T^{-1} \quad (1.1)$$

In the presence of the external magnetic field, $\vec{B}$, there is a torque on the magnetic moment, which changes the angular momentum vector and causes L and m to precess around the direction of $\vec{B}$ at a Larmor frequency of $\gamma_p \vec{B}$, shown in Figure 1.4 [2].

When the current (directed magnetic field) is removed from the coil, the protons precess back in alignment with the ambient geomagnetic field, producing a small alternating current in the coil at the Larmor frequency [31]. Knowing the Larmor frequency and magnitude of the

directed magnetic field enables accurate determination of the total-field intensity.

Advantages of proton precession magnetometers include their low cost and ease of operation for ground and aircraft-based surveys. No leveling or orientation adjustments are required of the sensor head beyond keeping it at a high angle with respect to the geomagnetic field, and the instruments are typically precise to 1 nT [40]. Proton precession magnetometers also lack moving parts, making their operation robust in a variety of environments and weather conditions [31]. The low power requirement for many models also enables continuous operation and rapid data collection in the field [2].

Disadvantages of proton precession magnetometers include the interpretability of their scalar outputs and potential sensor and electronics issues that may affect the accuracy of their measurements. Since they provide information about the total field intensity, not direction information, data are typically more difficult to interpret than vector magnetometers [40]. The accuracy of the proton precession magnetometer is impacted by (a) errors in frequency measurement, (b) distorted field measurements from magnetically unclean sensors, and (c) mechanical rotation, which can add or subtract to the precession frequency [2]. Additionally, since the measured precession data are small, it is important to operate proton precession magnetometers as far away from unwanted magnetic




Figure 1.4: Precession around an external magnetic field. A proton's angular momentum, L , and parallel magnetic field vector, $\vec{m}$, changes in the presence of an external magnetic field, $\vec{B}$, causing a precession around $\vec{B}$ proportional to the proton-gyromagnetic ratio. (Adapted from [2])




Figure 1.5: Fluxgate magnetometer diagram. Primary coils and oppositely-wound secondary coils surround two ferromagnetic cores. Driving current through the primary coils will induce an opposite magnetic field from the primary coil. When the cores are aligned with an external magnetic field, the resulting asymmetry in the secondary coil will produce a voltage proportional to the magnitude of the external field. (Source: [2])

noise sources such as AC power lines and cultural sources such as cars, buildings, and utility lines as possible.

Fluxgate Magnetometers Fluxgate magnetometers measure the vector components of a magnetic field by exploiting the nonlinear relationship between an external magnetic field and induced magnetization in ferrimagnetic materials [2, 41]. The basic fluxgate sensor typically consists of two ferromagnetic cores with coils of wire wound around them. Both cores are wound with a primary coil. Still, the orientation of the winding is opposite one another such that a drive current will saturate the cores and induce magnetic fields with equal magnitude but opposite direction, effectively canceling out the magnetization in the absence of an ambient magnetic field [31]. Secondary coils surround the primary coils and are oppositely wound in series. When no external magnetic field is present, driving current through the primary coils results in their magnetic saturation but no induced voltage in the secondary coils since the primary coil magnetizations were canceled. Figure 1.5 shows a diagram of the basic fluxgate magnetometer.

When the cores are aligned with an ambient magnetic field, an asymmetry arises between the magnetization of the cores, causing an asymmetry of the rates of change of magnetic flux in the secondary coils and inducing an electromotive force in the sense winding at twice the frequency of the driving current signal (i.e., the second harmonic) [2, 42]. The second harmonic produced is proportional to the component of the ambient field that is parallel to the orientation of the core axis. A triaxial fluxgate magnetometer may consist of two




Figure 1.6: Ring-Core Fluxgate Magnetometer Diagram. Source: [3]

individual ring-core fluxgate sensors whose sense windings each measure one of the orthogonal magnetic field components. One of the ring cores has windings perpendicular to one another, enabling the measurement of two of the field components simultaneously. The second ring core has windings oriented perpendicular to the first ring core, allowing for measurement of the third component. Figure 1.6 shows a schematic of a typical ring-core fluxgate sensor (Figure 1 of [3]).

The advantages of fluxgate magnetometers include their light weight, robust architecture, and low power consumption [43]. Most commercially available fluxgates are relatively low-cost, facilitating the construction of multi-sensor packages such as magnetic gradiometers. These can overcome some inherent instrument limitations and allow for easier calibration [44]. Their high sensitivity and high-resolution capabilities make fluxgate magnetometers particularly well-suited for shallow geophysics applications with low-metal targets [45]. Low power consumption

and lightweight fluxgate sensors allow higher ground cover efficiency during surveys, despite necessary calibration tasks [32].

Disadvantages of fluxgate magnetometers arise from measuring relative changes and include intrinsic errors in sensitivity, offset, and angle [43]. Since most fluxgates measure with 1 nT accuracy, the orientation of the sensor head must be known to be within 0.001° to ensure proper alignment, which may be cumbersome when used as a hand-held instrument in the field (this is discussed further in Chapter 3). Relatively large offset drift has been observed in triaxial fluxgate magnetometers, which may be reduced in either the analog or digital domain [44]. Most commercially available triaxial fluxgates include electronic feedback arrangements to linearize the measured response, so the use of three independent units near one another may result in feedback errors [2].

1.3.2.2 Ground Penetrating Radar (GPR)

GPR detects buried objects by sending radio waves into the ground and observing return signals generated by reflections off interfaces with contrasting electrical properties [32]. The GPR system consists of a signal generator, transmission antenna, receiving antenna, and data acquisition electronics [31]. Commercial units typically have a fixed geometry for the transmitting and receiving antennae, which are moved over the ground surface to detect subsurface features [46]. Since electromagnetic (EM) waves propagate with very little dispersion, the scattered or reflected signal is very similar in appearance to the transmitted signal, making reflected signal recognition straightforward [47]. Analyzing the time delay, signal intensity, and phase between transmitted and received signals supplies information about the distance, size, and velocity of reflectors in

the subsurface.

Transmission of the radio waves occurs as either a pulsed-echo (time pulse), frequency-modulated continuous wave (FMCW), or stepped frequency continuous wave (SFCW). In a pulsed-echo GPR, the radio wave is emitted in a repetitive train of short bursts around a central frequency, and a return signal is received between timed pulses [47]. The transmitted signal for a pulsed-echo system is typically a Ricker wavelet, the second derivative of a Gaussian function. FMCW GPR units emit continuous transmission power but can change operating frequencies during the measurement. This creates a sweep across which frequency and phase differences between the transmitted and received signals are collected [48]. SFCW GPR units emit discrete waves while continuously increasing the frequency by a constant step, collecting information about phase and frequency differences between transmitted and reflected signals at each discrete step [49].

The most common GPR antenna geometries are resistively loaded bowtie, log-periodic dipole, planar spiral, or Vivaldi TEM Horn [50,51]. Figure 1.7 shows the most popular designs from commercial GPR units and their 3-D radiation patterns.

An individual reflection is called an A-scan, and the combined A-scans over a survey transect form a B-scan, otherwise known as a radargram [52]. A volumetric representation of the survey area may be constructed with proper gridding and a sufficient number of tie points. In the presence of an anomalous zone, the reflection arrival time will decrease as the GPR unit approaches the zone, reaching a minimum once directly above it, and increasing when moving away, resulting in a diffraction hyperbola on the radargram whose move-out corresponds to the radar velocity of the propagation medium [31,47]. Figure 1.8 shows a diagram of a typical GPR survey setup, including an example A-scan, B-scan with a diffraction hyperbola caused by the




Figure 1.7: Common GPR antenna geometries and 3D radiation patterns (right side of each antenna). (a) Circular Plate Bow-Tie. (b) Triangular Plate Bow-Tie. (c) Horn. (d) Vivaldi. (e) Archimedean Spiral. (f) Equiangular Spiral.




Figure 1.8: A typical GPR survey setup. The lawnmower-style commercial GPR unit sends EM waves into the ground, producing pulse echoes (A-scans). As the GPR moves across a survey line, the accumulated A-scans combine into a B-scan, known as a radargram, showing the subsurface profile in 2-D. Surveying with sufficient tie points allows for the construction of a C-scan, the 3-D volume representation of the survey area.

presence of an AT landmine, and slice from the 3-D volume representation.

The smallest layer that can be resolved in a GPR survey has thickness $\lambda/4$, and λ (wavelength) is given by $\lambda = v/f$, where v is the radar velocity, and f is the frequency of the GPR antenna [31]. Radar velocity is controlled by the material's dielectric permittivity and magnetic permeability through which the wave is propagating. Typical minefield radar velocities vary from 0.05 m/ns (for organic-rich topsoil) to 0.15 m/ns (for dry sand), and are significantly influenced by soil moisture [53]. To resolve a landmine, which may be 6 – 30 cm in diameter and buried at < 0.5 m depths, ultrawideband (500 MHz-4.5 GHz) antenna frequencies are required. Processing

techniques for the collected GPR reflections are analogous to those in reflection seismology and include common midpoint stacking and migration, gain recovery, spiking deconvolution (wherein the minimum-phase equivalent portion of the embedded mixed-phase GPR wavelet is estimated from the reflection data and removed, leaving an all-pass filter that is pure phase-shift with unit amplitude spectrum), velocity analysis, bandpass filtering, NMO corrections, static and elevation corrections, and f-k migration [54,55].

Advantages of using GPR include the retrieval of accurate burial depths, an abundance of available processing software, and the ability to detect a variety of buried target materials [25]. When a reasonable radar velocity estimate is possible for a survey area, GPR can assist in determining reflector burial depth with high ($< 95\%$) accuracy in most circumstances [25]. Since data processing steps are similar to well-established seismic methods, many proprietary and open-source software suites are available to process and interpret GPR data [56]. Most landmine and ERW casings provide high dielectric contrast to the host overburden, so GPR can see reflections from the majority of casing materials, depending on size, shape, and burial depth [52].

Disadvantages of using GPR include suboptimal performance in certain environmental conditions, dependency on high dielectric contrasts, bulky instrumentation, and frequency limitations of commercially available systems [1]. Radar signals attenuate with increasing water content in the subsurface, making them unsuitable in high-humidity minefields such as Cambodia and the Falklands [56]. GPR systems rely on a sufficiently large dielectric contrast between buried targets and the host overburden to produce diffraction hyperbolas, so some targets may be more difficult to distinguish based on metal content, operating frequency, soil moisture values, and the roughness of the ground surface boundary [32]. GPR systems are often large and heavy, depending on antenna frequency, which may be cumbersome to transport and operate in a minefield [25].

1.3.2.3 Visible Light and Infrared (IR)

Visible light methods detect buried objects by capturing images of a survey area and exploiting the smoothness of the object when compared to the natural environment. The main uses of visible light for landmine detection include color and intensity analysis, polarization analysis, and edge detection and grouping [57]. In color and intensity analysis, the global statistics of an image captured with a visible light camera are used to set a contrast threshold above which an object is identified as a potential target [58]. Polarization analysis involves comparing the polarization of an image with its background to detect anomalous objects using either a specialized polarization camera or visible light camera fitted with a polarization filter [59]. Edge detection and grouping analyze visible light images, searching for straight or circular edges using Hough transforms to identify artificial objects from the surrounding natural environment [60].




Figure 1.9: Example of thermal (IR) imagery used to detect buried plastic and metal mines. (Adapted from Figure 6 of [4])

IR imaging detects buried objects and subsurface anomalies by measuring the variations in electromagnetic radiation emitted or reflected from the objects and surrounding soil or the objects and their overlying vegetation [32]. Since landmines and other ERW release or absorb heat at a different rate than their surrounding host material, a thermal contrast will be created as either a volume effect (resulting from the disturbance of heat flow) or a surface effect (resulting from the disturbance of

the topsoil during burial) [52]. Volume effect thermal disturbances persist long after initial landmine burial thanks to diurnal variations, though the magnitude of thermal differences will decrease with increasing burial depth [60]. Commercially available IR cameras operate at either short wavelengths (0.9-1.7 μm), mid-wavelengths (2 – 5 μm), or long wavelengths (7 – 12 μm). Figure 1.9 shows the positive detection of buried metal and plastic mines with IR imagery taken from nadir (90°) look angle. To detect landmines buried up to 15 cm depth, 8 – 12 μm wavelengths are required [61]. IR cameras can be either active (where IR LEDs are used to cast light over the area before image capture) or passive (where thermal radiation is detected without the need for additional illumination) [62]. IR and visible light imagery may be combined to increase surface-laid target detection potential, as demonstrated in [63–65].

Advantages of using visible light and IR include the ability to detect mines long after their initial burial, lightweight instrumentation, and the ability to detect a variety of casing materials [66]. While surface thermal effects will degrade within weeks of initial burial, volume thermal effects and reflectivity contrasts remain long after, enabling the use of visible light and IR in older minefields as well as active warzones [60]. The majority of visible light and IR cameras weigh less than 0.5 kg, allowing for ease of integration on handheld, ground-based, or aerial detection systems [67]. Since visible light and IR methods exploit the reflective and thermal differences of landmine casings from their surroundings, both metal and non-metal targets can be detected [68].

Disadvantages of using visible light and IR include the requirement of sensitive cameras, strong dependence on environmental conditions, and detection limitations based on target material, orientation, and burial depth [66]. To achieve sufficient sensitivity and spatial resolution for ERW detection, IR cameras require the detection of $\Delta T < 0.1^\circ\text{C}$, which limits the use of many commercially available instruments [52]. However, visible light camera requirements are

far less limiting, including those with polarization capabilities. Since they rely on the thermal or reflectivity contrasts of buried objects with their surroundings, the presence of dense vegetation, low thermal contrast times (morning and evening), sun position, and the presence of clouds will significantly impede the accuracy of each technique [68]. Inhomogeneities in the subsurface and voids or clutter may lead to high false detection rates. Detection of artificial objects relies on the presence of straight or curved edges, so the burial orientation of a landmine or ERW will impact the efficacy of visible light, in particular [1]. Additionally, objects buried deeper than 15 cm may not be detected by visible light *or* IR methods.

1.3.2.4 Light Detection and Ranging (LiDAR)

LiDAR works in both infrared and visible regions of the EM spectrum to detect objects by sending out bursts of coherent radiation and measuring reflections from objects at the surface [35]. Traditional LiDAR systems calculate the distances between the instrument and a detected object by comparing the transmitted and reflected/absorbed energies, while polarization LiDAR systems search for the smoothness of objects compared to their surroundings, providing additional information about artificial vs naturally occurring material [1]. Data collected from LiDAR produce ground surface contour maps that reveal small-scale ground disturbances potentially associated with munitions-related activities [69]. Figure 1.10 shows the results of using UAV-mounted LiDAR to map minefield features in Angola [5]. Since LiDAR is unable to detect buried objects, it is often used in tandem with other geophysical techniques for landmine detection [70, 71].

Advantages of LiDAR include operational safety, metal and non-metal detection capabilities, and robust performance in a variety of climate conditions. Modern LiDAR systems operate




Figure 1.10: Example of using LiDAR to map minefield features. (Adapted from Figure 7 of [5])

at either 905 nm or 1550 nm, which are safe to work with for UAV-based minefield surveying as safety concerns arise only when looking directly at the light beam [72]. Since metal and non-metal landmine and ERW casings will have different reflection/absorption properties than surrounding soil, LiDAR can detect various surface-laid targets. LiDAR operates in both day-light and night conditions, as well as in the rain, as the sensor is below the cloud line [1].

Disadvantages of LiDAR include its inability to detect buried objects, suboptimal performance in densely vegetated areas, and limitations of range and sensitivity. Since LiDAR uses shorter wavelengths than GPR or IR techniques, signals are unable to penetrate the subsurface, limiting detectable targets to those laid on the ground surface or with at least a portion of the casing exposed at the surface [73]. Surface obstructions such as dense vegetation will impede LiDAR use and accuracy, as the vegetation results in low return signal strength [1]. Additionally, range and sensitivity limitations arise when operated in areas with rapidly changing landscapes,

which create clutter in returned signals.

1.3.3 Multi-Sensor Systems

Commercially available multi-sensor systems consist of handheld dual-sensor systems or a single-sensor type array. Dual sensor systems such as the Vallon “Minehound” consist of a handheld metal detector combined with GPR [74]. While some of the environmental limitations of GPR technology are overcome by integrating a metal detector with GPR, the instrument still relies on the presence of at least a fraction of metal in the UXO casing for detection and puts the operator < 1 m from the blast radius, creating a potentially lethal situation for de-miners. Other dual sensor systems rely on combining gradiometer and magnetometer technology, with similar environmental limitations and proximity issues for end users. Instruments such as GapEOD’s UltraTEM system combine an array of transient electromagnetic sensors on a ground-based vehicle to detect UXO, and require a significant amount of metal to be present to be effective, which is not common for many varieties of mines [75]. Geometrics MetalMapper is a ground-vehicle towed time domain electromagnetic system with three orthogonal transmitter loops and seven tri-axial receiver loops. The system uses electromagnetic induction (EMI) to detect buried metal objects and cannot detect non-metal or low-metal UXO [76]. Other EMI systems include the BUD “Berkeley UXO Detector,” which is an active, cart-based, electromagnetic induction sensor system that uses three orthogonal transmitters and eight pairs of differenced receivers to detect and estimate the geometry of induced currents from UXO [77, 78]. The BUD and MetalMapper are multi-sensor arrays of single sensor type.

Other ground-vehicle towed systems include TEMDAS, which is a time-domain EMI sys-

tem with 25 co-located transmit and receive pairs arranged in a 5x5 grid, designed to discriminate metal UXO from surrounding metal debris [79]. The GG Sciences Man-Portable Vector (MPV) is a mono/multi-static ultrawideband backpack and ground-based pushcart time-domain EMI system that measures three components of the secondary field at two different heights to detect metal UXO and estimate the burial depth [80]. The performance of the system is limited by its bulky size and high power consumption, allowing only short-duration surveys, impractical for use in rural minefields. Despite their advancements to the state of the art, all currently available ground-based multi-sensor systems are limited by steep terrain and unreliable weather conditions in minefield areas.

1.4 Outline of Thesis

In Chapter 2, we present our equipment and sensor specifications and an overview of our general approach to the landmine, ERW, and UXO detection problem. We evaluate the potential use of ground-based GPR for detecting targets at a realistic inert minefield in Pawnee, OK, and discuss the challenges and limitations of adapting commercial off-the-shelf (COTS) sensors for UAVs. We detail the COTS sensors used in our studies (GSSI 400 MHz GPR, FLIR Duo Pro R Camera, GeoSLAM Zeb Horizon LiDAR) and describe our design and fabrication process using MEDA FVM400 fluxgate magnetometer sensor heads as the base for an in-house built magnetic gradiometer system that we call TetraMag. We include schematics of our overall methodology and signal processing techniques utilized when analyzing our data. Finally, we describe the challenges of sourcing an appropriate drone system to meet our operational and research-focused needs, then describe the drone platform we integrated our sensor payload on (Freefly Alta X).

Chapter 3 presents results from a computational and field study of magnetometry for detecting small anomalies like landmines and UXO using finite-difference magnetic gradiometry (FDMG). Magnetometry is used to detect ferrous objects at various scales, but detecting small-scale, small-amplitude targets in the shallow subsurface remains challenging. Magnetic anomalies are often approximated as dipoles or volumes of dipoles that can be located, and their source parameters are characterized using scalar or vector magnetometers. Both types of magnetometers are sensitive to space weather and cultural noise sources that map temporal variations into spatial variations across a survey area. Vector magnetometers provide more information about detected bodies at the cost of extreme sensitivity to heading errors that cannot be reliably measured in the field. Magnetic gradiometry solves the problem of temporal-to-spatial mapping. It eliminates distant noise sources, but the heading error challenges remain, motivating the need for full magnetic gradient tensor (MGT) invariants that are relatively stable to rotation. Here, we show that the finite size of magnetic gradiometers compared to the lengthscales of magnetic anomalies due to small buried objects affects the properties of the gradient tensor, including its symmetry and invariants. This renders traditional assumptions of magnetic gradiometry largely inappropriate for detecting and characterizing small-scale anomalies. We present our developed methodology that leverages properties of the finite-difference MGT and its invariants appropriate for small-scale, small-amplitude targets in the shallow critical zone, like UXO and landmines. We validate these results with synthetic and field data using a triaxial magnetic gradiometer (TetraMag).

In Chapter 4, we present the result from a study of machine learning (ML) methods for determining magnetic source parameters of targets from FDMGT invariants and detecting the presence of UXO in GPR radargrams. We describe an analog UXO field site where we collected GPR data, then detail a convolutional neural network (CNN) designed to extract features from the

resulting radargram images, quantifying the benefits and limitations of the technique. We find that testing accuracies range from 88-95%, suggesting a CNN is an appropriate tool for identifying radargram transects of interest in a UXO-contaminated area. We then extend the use of CNN to magnetic gradiometry data using synthetic FDMGT invariant maps as inputs to the network. We combine different invariants as input layers to the CNN, adding a linear regression layer to the end of the network for estimating the magnetic moment and burial depth of a detected ferrous target. We find the CNN can estimate magnetic moment within 0.6 Am^2 at 99% accuracy, amounting to an order of magnitude improvement compared to current methods [81, 82]. The CNN is also able to estimate depth within 4 cm at 99.43% accuracy, amounting to nearly an order of magnitude improvement on current depth estimation techniques for small targets [83].

Chapter 5 presents a method of modeling and reducing unwanted magnetic signals generated by vehicle self-noise. Drone-based magnetic surveys have a wide range of applications in geoscience and have gained popularity due to the increasing availability and decreasing cost of commercial drones. Challenges for drone-mounted magnetometry include the generation of unwanted magnetic self-noise from DC motors and power systems, which can obfuscate small-scale, small-amplitude anomalies like UXO and landmines. While slung-loaded magnetometer systems are used to avoid these unwanted signals, they are best suited for total field magnetometers. They may be inappropriate for surveys requiring very low-altitude flying. To mitigate space weather and cultural noise sources, and maximize the amount of information collected in the field, we utilize magnetic gradiometry with an array of rigidly-mounted triaxial fluxgate sensors (TetraMag), aboard a Freefly Alta X drone. This necessitates characterizing the combined magnetic signature of the drone to remove from measured data. Utilizing telemetry data from the Alta X at an inert minefield in Glenn Dale, MD, we built a model for the expected magnetic field

based on the real-time kinematic (RTK) position, instantaneous current drawn from each battery to each motor, revolutions-per-minute (RPM) from each motor, and the location of each magnetometer sensor head. We then used this model to subtract the expected magnetic field strengths, component-by-component, from the measured data at each sensor location. We find multiple orders of magnitude noise reduction with our approach. Applying this model to collected data improves detection fidelity and enables higher-resolution localization of detected targets, which we show using FDMGT invariant maps. Many commercial drone systems collect these telemetry data by default, so leveraging them to characterize the combined self-noise magnetic signature of the vehicle is an effective and efficient way to enable the use of rigidly-mounted vector magnetometers on drones, especially for systems where slung-loaded solutions are inappropriate.

Chapter 6 presents the results of a full multi-modal geophysical study of the inert minefield in Glenn Dale, MD, using our payload-equipped UAS. We combine methods described in previous chapters to assemble FDMGT invariant maps of the field area, integrating them with 3-D LiDAR, visible light photogrammetry, and thermal data to fully describe and characterize the targets at the field site. We describe the efficacy and practicality of each sensor modality in terms of detection potential and the benefits and limitations of combining the datasets. We find that the combination of sensor modalities and methods developed here facilitates a more complete picture of the geophysical parameters of a field site and may be used not only for the detection of dangerous anthropogenic sources like landmines and UXO but also for broader shallow critical zone geophysics studies.

Chapter 7 provides the conclusion to the thesis, detailing lessons learned from this work and ideas for future research that build on our results.

Chapter 2: Equipment, Sensors, and General Methodology

In this chapter, we describe and discuss our geophysical sensor selections, the development of necessary equipment to achieve the goals of this work, and the integration of our payload package on a UAV. We then summarize our general approach and detail our system's communications and controls setup.

2.1 Geophysical Sensors

2.1.1 GPR Technical Specifications

The GPR used for these studies is the GSSI SIR 4000 with a model 50400S 400 MHz shielded analog antenna mounted to a plastic push-cart. The data acquisition portion of the system is single-channel, weighing 4.5 kg with 32 GB Flash storage capacity and 1 GB RAM. The model 50400S antenna has a center frequency of 400 MHz (200-600 MHz range), which translates to an average depth range of 0-4 m, depending on the radar velocity of the medium through which the radio waves are propagating. Figure 2.1 shows the GPR setup. The antenna is housed in the orange case, and users can see real-time data visualization using the tablet connected to the push bar.

Data from the GSSI GPR are processed using proprietary GSSI software (Radan7) and



Figure 2.1: GSSI SIR 4000 pushcart with 400 MHz antenna.

open-source (MatGPR, written for MATLAB). All GPR data collected for these studies were processed using the following standard steps:

- Time Zero Correction: adjusts the zero time position of the air and ground waves from the GPR wavelet, ensuring accuracy of the calculated depth profiles.
- Background Signal Removal: reduces and/or removes background clutter from the signal to improve the signal-to-noise ratio.
- Filtering: noise removal and fine-tuning of the signal to ensure optimal interpretability (asymmetric filters were avoided).
- Gaining: magnifies desired sections of the radargram, highlighting areas of interest (using Standard AGC).
- Migration: corrects for the geometry and spatial distribution of the survey area for depth migration using a 1-D velocity model constructed using MatGPR Model Builder.



Figure 2.2: Condensed target grid at humanitarian minefield in Pawnee, OK. (Figure 3 from [6])

We conducted two ground-based studies, detailed in the following sections, to evaluate the practicality of using GPR for landmine, UXO, and ERW detection.

2.1.1.1 Evaluating Ground-Based GPR in Realistic Minefield Conditions

Collaborators from the Demining Research Community (DRC) constructed an inert minefield at the University of Oklahoma's Center for Fire and Explosives (CENFEX) facility in Pawnee, Oklahoma [6]. The study site is in a soil map unit named for the Renfrow soil series (Fine, mixed, superactive, thermic Udertic Paleustolls), with rough terrain representative of realistic minefield conditions. The 10 m x 40 m grid contained 143 seeded targets, including control holes, clutter, landmines, and UXO. Figure 2.2 shows a condensed grid visualization for the inert minefield constructed by DRC (Figure 3 of [6]). Each row was spaced 2 m apart from one another, and each target within a row was 1.5 m apart.

The GPR survey consisted of 208 total scans with 50 cm spacing, which we conducted

in N-S and E-W orientations. The survey took place in hot and humid conditions after a week of consistent rain, which is suboptimal for GPR as attenuation increases with increasing soil moisture. Rough, uneven terrain and the presence of surface-laid targets resulted in significant localization errors, as the cart needed to be lifted to overcome local obstacles, not allowing the scroll wheel to keep track of the distance accurately traveled along each transect. The data were difficult to process and interpret due to poor quality.

While clear hyperbolas were visible in some of the processed radargrams, localization errors of 1 m or more would be unacceptable for a de-miner to use in the field. Figure 2.3 shows two example radargrams from the GPR dataset. The top radargram contains some clear hyperbolas, indicated by a dashed circle, and localization for these is well-constrained since the scroll wheel maintained contact with the ground for the whole transect. However, hyperbolas in the bottom radargram can not be localized well due to multiple instances where the instrument needed to be lifted (indicated by black arrows). Figure 2.4 shows an overlay of full or partial hyperbolas picked from each 2-D radargram (black dots) with expected target locations (red stars). Hyperbola locations reflect only the latitude and longitude of the hyperbola vertex, as depth estimates cannot be reliably constrained based on the dataset. While some of the hyperbolas cluster near expected target locations, the uncertainty of the GPS location and the uncertainty in scroll wheel position along each transect make it difficult to distinguish any single target in the field area. Data quality issues and excessive operational challenges for ground-based surveying motivate exploring the potential for GPR use on a UAV instead.




Figure 2.3: Example processed radargrams from GPR inert minefield study. The top radargram has clear hyperbolas visible (indicated by dashed circles) and good localization. The bottom radargram shows some hyperbolas, but localization is poor due to the GPR being lifted multiple times during the run (indicated by black arrows).




Figure 2.4: Expected target and GPR hyperbola locations.

Reflector	Diameter	Burial Depth
 AT Landmine	24 cm	13 cm
 AP Landmine	9 cm	13 cm
 Analog UXO	15 cm (length)	12 cm

Figure 2.5: Summary of target types, their dimensions, and burial depths used for height-varying GPR experiment.

2.1.1.2 Evaluating Commercial GPR for UAV Integration

We conducted a series of height-varying experiments to evaluate the potential for the COTS GSSI GPR to identify landmine and UXO targets from a UAV at operational flying heights. Figure 2.5 summarizes the type of targets, their dimensions, and their burial depths. The targets were buried 1 m apart along a single transect. The analog UXO, a ferrous metal water bottle, approximates large artillery shells' dimensions and material properties. The FCC limits operational flying heights to < 1 m for GPR frequencies [84], so we restricted our maximum testing height to 81 cm. We used the GSSI 400 MHz GPR on a sliding cart to scan the transect from three heights: (1) 5 cm, (2) 44 cm, and (3) 81 cm. The data were processed using the steps outlined in Section 2.1.1: a full pass filter for background noise removal, FIR filter for anti-aliasing, AGC for gain, a 200-700 MHz band pass filter, and Kirchoff depth migration.

Figure 2.6 shows the processed radargrams from the height-varying survey. The locations of the AT mine, AP mine, and analog UXO (ferrous metal water bottle) are indicated by arrows. The top plot shows the radargram for an antenna height of 5 cm off the ground. Hyperbolas are visible for all three targets (circled with a solid line), and the strongest signal comes from the AT landmine. The middle plot shows the radargram when the GPR antenna is 43.6 cm off the ground. Significant side reflections are visible as wavy patterns in the radargram, though hyperbolas are partially visible for all three targets (circled with a dashed line). For an operator analyzing these data, however, it is unlikely those hyperbolas would be identified. By 81 cm off the ground (bottom plot), the side reflections and power loss are so severe that none of the targets are detected. Side reflection clutter increasing with height is likely a product of the antenna shape used. The model 50400S antenna is a bowtie shape, which produces a donut-like radiation pattern (seen in Figure 1.7). When this antenna shape is ground-coupled with an upper reflection plate, the majority of the energy will be directed into the ground. As the distance between the ground and the antenna increases, the lack of directivity results in significant power loss as the radiated signal leaks from the side lobes. The challenges of working with the ground-based GPR in realistic minefield conditions and the power loss at very low altitudes resulted in our decision not to include the GSSI GPR system in the UAS integration.

2.1.2 Dual IR and Visible Light Photogrammetry

As described in Chapter 1, IR (thermal) and visible light photogrammetry are useful for detecting landmines, UXO, and ERW. For this work, we chose the FLIR Duo Pro R, an 85 x 81.3 x 68.5 mm dual IR and visible light camera, weighing only 325 g. The thermal imager operates




Figure 2.6: Radargrams from height-varying GPR survey over landmines and analog UXO.



Figure 2.7: FLIR Duo Pro R dual visible light and thermal camera.

in the 7.5-13.5 μm spectral band, appropriate for up to 15 cm depth imaging, with 640 x 512 resolution. The thermal lens is 19 mm, providing 32° x 26° field of view at 30 Hz frame rate. The visible sensor has 4000 x 3000 resolution and a 56° x 45° field of view. The low-power requirements and integrated GPS make this an ideal choice for UAV-based surveying. Figure 2.7 shows the FLIR Duo Pro R. All visible light and thermal image data were processed using scripts written in MATLAB and Python.

2.1.3 LiDAR

LiDAR is useful for draping other sensor modality data on and examining textural features from return intensity that may indicate the presence of landmines, UXO, and ERW. For this work, we chose the GeoSLAM Zeb Horizon 3-D LiDAR, a relatively lightweight (< 3 kg) unit that uses a Class 1 (903 nm wavelength) laser to produce a 3-D point cloud. The Zeb Horizon is capable of scanning 300,000 points/sec using its 16 sensors with a relative accuracy of 6 mm. The vertical angular resolution is 2°, and the horizontal angular resolution is 0.2°, making it one of the most competitive COTS LiDAR units available. The unit has integrated simultaneous localization



Figure 2.8: GeoSLAM Zeb Horizon 3-D LiDAR.

and mapping (SLAM) software built-in, allowing the LiDAR to track its position as it moves in space accurately and allowing for ease of use for both handheld walking surveys and UAV-based surveys. Figure 2.8 shows an image of the GeoSLAM Zeb Horizon.

Initial data processing and point cloud registration for collected data rely on GeoSLAM proprietary software to produce point cloud files. MATLAB and CloudCompare, an open-source 3-D point cloud and mesh processing software, are used for further analysis of the data.

2.1.4 TetraMag

As we discussed in Chapter 1, fluxgate magnetometers provide valuable component information that can be used to detect and characterize small anomalies but suffer from orientation and calibration errors when used in isolation. To maximize the magnetic data collected and overcome some of the practical limitations of a single fluxgate magnetometer, we designed and fabricated a triaxial fluxgate magnetic gradiometer called TetraMag. Gradiometers are sensitive to nearby changes but insensitive to distant changes, solving the temporal-to-spatial mapping of cultural and space weather noise. Having redundancy of tensor components enables inherent error correction and noise estimation. Version 1 of our design was primarily used for proof-of-concept and informed the final design choices for Version 2, which we integrated with our UAV. MEDA

FVM400 triaxial fluxgate sensors were used as the base sensor, and we set up a Raspberry Pi microcontroller fitted with a real-time clock (RTC) for collecting and storing data output from the FVM 400 sensors using Python scripts and wireless communications. We ingested, parsed, and analyzed all TetraMag data using MATLAB scripts.

2.1.4.1 TetraMag v1

Based on the work of Xu et al. [85], we selected a regular tetrahedron shape with four MEDA FVM triaxial fluxgate sensors placed at the vertices. A tetrahedron is a mechanically stable geometry choice because stresses are transmitted primarily along the edges, so strains become mostly compressional or extensional and not flexural. This allows us to sample the full magnetic gradient tensor in a finite-difference sense, which is discussed more in Chapter 3. The sensor heads were placed in the same orientation at each vertex, minimizing the relative temporal angular variations from thermal expansion or mechanical strain. Figure 2.9(a) shows the CAD rendering of our initial design. The center hub and sensor housings were 3-D printed using PLA filament. Carbon fiber tubes connect the sensor housings to the center hub, and an aluminum plate connects to nonmagnetic 80/20 aluminum extrusion support beams using brass hardware. We selected the scale of TetraMag to be 1 m (side length 0.5 m) since the anomalies will be measured from 1 m flying height. We integrated two 1-D laser rangefinders pointing orthogonally from one another for spatial localization.

To test the capabilities of TetraMag at various heights at the UMD Paint Branch Turfgrass research facility, we constructed a nonmagnetic scaffold system, seen in Figure 2.10. The aluminum top plate from TetraMag v1 is connected to an 80/20 slider and rail that can be adjusted




Figure 2.9: TetraMag initial design (a) and with added structural supports (b).

by lowering or raising the support rail on the rungs of the scaffold. We designed 3-D printed coupling pieces that attach the legs of the scaffold to MRI-safe nonmagnetic wheels that allow the whole scaffold to roll on fiberglass U-channels. We simulated lawnmower-style flight paths by rolling the scaffold and sliding TetraMag along its rail, scanning over three targets: (1) AT landmine, (2) AP landmine, (3) a ferrous metallic water bottle of approximate size and material properties as a medium-scale UXO. We found during initial tests that structural modes of TetraMag were being excited by the torsional motion of the bottom three sensors when moving. Since the frequency of these modes lies within the Nyquist frequency for the system (2 Hz), we designed an additional support system to limit the torsion of the bottom sensors and put the center hub under tension, seen in Figure 2.9(b). This increased the structural modes beyond the Nyquist frequency and enabled us to fine-tune the gradiometry techniques described further in Chapter 3.




Figure 2.10: TetraMag field test setup using a height-adjustable nonmagnetic aluminum scaffold.



Figure 2.11: TetraMag v2, designed as landing gear for the Freefly Alta X UAV.

2.1.4.2 TetraMag v2

We designed a second version of TetraMag for direct integration with a Freefly Alta X UAV. To reduce the weight of aluminum support plates and beams and to avoid low-frequency structural modes, we designed TetraMag v2 as rigid landing gear for the drone. We maintain a regular tetrahedron shape with 0.5 m side length by embedding the tetrahedron in an octohedron-like frame, which acts as a safety buffer to protect the fluxgate sensors in the event of a crash. We designed and 3-D printed all connecting pieces using NylonX filament to provide higher rigidity and temperature tolerance. Carbon fiber tubes connect the sensor housings and edge pieces, which were then epoxied in place. Figure 2.11 shows a CAD rendering of the TetraMag v2 design. The top plate is made from carbon fiber and matches the bottom shape of the Freefly Alta X.

2.2 UAS Integration

This work required a high-payload capacity drone capable of carrying our equipment. We integrated our sensors on a Freefly Alta X, a LiPo battery-powered quadcopter UAV with a 15 kg payload capacity. Figure 2.12 shows the expected flying time by payload weight [7]. The Alta X maintains high-precision localization using RTK GNSS, and altitude measurements for flights < 5 m are made with a laser range altimeter we integrated into the vehicle. The UAV is controlled via two systems: a ground station laptop running QGroundControl software, and a handheld Futaba T14SG RC controller. All flight plans used QGroundControl and lat/lon positions as waypoints for the UAV to fly to during a mission. The Futaba RC controller is used for non-mission flying and as a backup in case the mission being flown becomes unstable or dangerous. Figure 2.13 shows a CAD rendering of the Alta X with our custom TetraMag landing gear (top) and a photo of the complete UAS. Red arrows indicate sensor locations.

2.3 General Approach

Figure 2.14 summarizes our approach to detecting small-scale, small-amplitude targets such as landmines, UXOs, and ERWs. While this is our particular use case, it should be noted that these geophysical modalities and our general approach may be applied to larger-scale geophysical exploration, underground utility mapping, and general geoscience investigations of the shallow subsurface, much of which overlaps with our submitted patent application [86]. We couple geophysical sensors that measure thermal, visible light, magnetic gradiometry, and LiDAR data with the UAV's telemetry sensors that measure motor RPM, motor current draw, temper-

Alta X Approximate Flight Time vs Payload for 2 x 16Ah Batteries




Figure 2.12: Alta X flying time vs payload weight. Source: [7].

ature, and humidity to detect targets with various casing materials. We use LiDAR to provide a base map on which we drape other data and to find textural features that can indicate void spaces, recently disturbed ground, or reveal surface-laid targets. We use FLIR Duo Pro for visible light photogrammetry to detect surface-laid targets and thermal photogrammetry to detect objects from surface level to a maximum depth of 15 cm. We use TetraMag to detect small-scale, small-amplitude ferrous anomalies. Telemetry data are used to model the magnetic noise of the vehicle for noise reduction and high-accuracy positioning that collocates all sensor data.

Figure 2.15 is a diagram showing communications, controls, power distribution aboard the UAV, and how each sensor collects and stores data. Items with matching communication method symbols (Wi-Fi, Bluetooth, Radio Communication) exchange information with one another during data collection. Boxes with the SD card symbol represent locations where sensor data are stored and offloaded for post-flight processing.



Figure 2.13: Freefly Alta X with integrated drone design (top) and implementation (bottom).




Figure 2.14: General approach summary. Geophysical sensors (FLIR Duo Pro R, GeoSLAM Zeb Horizon, TetraMag) are coupled with telemetry sensors (measuring motor RPM, motor current draw, temperature, humidity) on a quadcopter UAV (Freefly Alta X).




Figure 2.15: Sensor Modalities, Controls, and Communications.

Chapter 3: Enabling Small Anomaly Detection using Finite-Difference Magnetic Gradiometry

3.1 Introduction

This section presents results from a computational and field study of magnetometry for detecting small anomalies like landmines and UXO using finite-difference magnetic gradiometry (FDMG). We discuss the challenges and limitations of current methods to detect and characterize small-scale, small-amplitude anomalies. We show that the finite size of magnetic gradiometers compared to the lengthscales of our magnetic anomalies of interest affects key properties of the gradient tensor and renders traditional assumptions of magnetic gradiometry largely inappropriate for our use case. We then present our developed methodology that leverages properties of the FDMG tensor appropriate for small targets in the critical zone. Finally, we validate our results using synthetic and field data collected with our TetraMag system.

3.1.1 Magnetometry in Geoscience

Magnetometry is a ubiquitous geophysical technique used to detect targets such as ore bodies [87–90], igneous intrusions [91–94], subsurface structures [95–98], and, for some humanitarian geophysics applications, landmines and UXO [99–101]. Magnetic anomalies are of-

ten approximated as dipoles or volumes of dipoles [102, 103], and localizing and characterizing their source parameters is challenging, whether using data from scalar or vector magnetometers. Geologic-scale features generally span spatial dimensions of 100s of m to 100s of km and produce anomalies on the order of 10s of nT to 1,000s of nT [104]. Landmines and UXO typically span spatial dimensions of 0.05 m to 2 m and produce anomalies on the order of 10s of nT to 100s of nT [37, 105].

Surveys to detect geologic-scale magnetic anomalies typically utilize scalar magnetometers, often performed aurally. Surveys to detect small-scale targets such as landmines and UXO utilize scalar and vector magnetometers and are often ground-based. However, efforts to integrate magnetometers on UAVs for landmine and UXO detection have gained traction over the past decade [62, 65, 106–108].

In this study, we focus on small-scale buried objects in the shallow critical zone ranging from infrastructure to landmines, which necessitates the characterization of small-scale, small amplitude signals. This task comes with unique challenges for survey methodology and instrumentation selection, which affect the fidelity and ability to interpret collected data. We investigate the impact of finite-difference vs. infinitesimal magnetic gradient tensors in the context of this challenge and consider how it relates to the size of the gradiometer, spatial distance, and azimuth from the target. We explore the utility of several magnetic gradient tensor invariants for target localization and then compare their performance under rotation using synthetic data. Finally, we validate the developed methods using field data for a small, dipole-like target buried at 15 cm depth.

3.1.2 Challenges of Scalar, Vector, and Gradient Magnetometry

3.1.2.1 Scalar Magnetometry

Scalar magnetometers are the standard tool used for ground-based and aerial surveys due to their high sensitivity and ease of data analysis. Crucially, scalar measurements do not depend on instrument orientation. This facilitates observations when precise pointing control is difficult to achieve, which is the case in typical field conditions, especially when flying. Their high sensitivity enables scalar magnetometers to easily record temporal variations in magnetic field strength caused by space weather and cultural noise sources. These temporal variations will map into spurious spatial variations in a magnetic survey where magnetic field strength measurements are made at different locations.

The magnitude of space weather contributions can range from pT to $100\ nT$ scales under geomagnetically quiet conditions, but variations of $1000s\ of\ nT$ can occur within minutes during geomagnetic storms [109,]. Proximity to urban magnetic noise sources such as a direct current (DC) metro system also significantly impacts the ability to reliably map out anomalies at the $100\ nT$ scale [110]. Temporal variations from either source can result in spurious spatial signals that swamp out the expected amplitudes from small targets of interest. When mapping large-scale anomalies, these spurious signals can be partially corrected by applying diurnal corrections with a base station and/or subtracting out the expected field for the survey location from the IGRF model [111–113].

While the precise mapping of temporal-to-spatial variations depends on the acquisition geometry, in general, short-period variations map to small-scale structures, while long-period

variations can map into either small or large-scale structures. Base station recordings are often utilized to correct spurious signals from total field and vector component magnetic data, especially in aeromagnetic and UAV-based surveys. [114] examined the total field coherence between base stations at varying distances and found that periods longer than 10 seconds have coherence of nearly 1 at short distances but are nearly incoherent at short periods. Additionally, using base station measurements from reference stations over 200 *km* away results in coherence of less than 0.5, limiting the efficacy of publicly available magnetic reference station data unless the survey area is close to the reference station and corrections are only needed for long-period variations [114]. It is important to note that these studies were conducted during geomagnetically quiet conditions.

Figure 3.1 shows the difference in coherence in geomagnetically quiet conditions and geomagnetic storm conditions across two pairs of magnetic reference stations: BOU – TUC (Boulder, CO, and Tucson, AZ), and FRD - BSL (Fredericksburg, VA, and NASA Stennis Space Center, MS). For the BOU – TUC pair, (a) shows geomagnetically quiet conditions with Kp index = 0, and (b) shows a G3 geomagnetic storm with Kp index = 7–. Coherence values are nearly 1 at low frequencies but fall off for all components around 50 *mHz* during quiet conditions, suggesting that using base-station data for corrections is suboptimal when short-period variations (small structures) are expected. During a geomagnetic storm, however, there is enhanced coherence for North and East components due to the sheet-like distribution of currents in the ionosphere, suggesting that the use of fluxgate triaxial magnetometers may be preferable during these conditions. Extreme orientation sensitivities make the use of fluxgate magnetometers prohibitively challenging, which is discussed in more detail in the next section.

Even with high-sensitivity magnetometers, sensor height is limiting when targeting small-




Figure 3.1: Coherence between magnetic reference stations during geomagnetically quiet conditions, Kp index 0 (a,c), and a G3 geomagnetic storm, Kp index 7 (b,d), at locations across the United States: Boulder, CO (a), Tucson, AZ (b), Fredericksburg, VZ (c), and NASA Stennis Space Center, MS (d). In quiet conditions, all components' coherence is high at low frequencies and falls off rapidly at frequencies > 50 mHz. Vertical and total field coherence falls off during geomagnetic storms, such as in quiet conditions, but North and East components show enhanced coherence at higher frequencies.




Figure 3.2: Expected magnetic total field anomaly (B_t) of a dipole-like anti-tank (AT) mine buried at 15 *cm* depth measured from various sensor heights, compared to typical amplitudes of geomagnetic storms and nearby DC Metro systems (shaded areas) and the noise floor of magnetometer used in this study, MEDA FVM400 (dashed line). At heights $> 0.5m$, geomagnetic storms and DC Metro signals can obfuscate the AT mine, producing false flags.

amplitude, shallow anomalies such as landmines and UXO. Figure 3.2 shows the expected total field magnetic anomaly for a dipole-like AT landmine buried at 15 *cm* depth, as measured from various sensor heights. As the magnetometer height increases, signals from geomagnetic storms and/or nearby DC metro systems completely obfuscate the amplitude of the AT mine, increasing the chance of false flags. AP landmines and cluster munitions of a similar size will have even smaller signatures. This is of particular importance for UAV-based surveying and surveys carried out near urban environments with DC metro systems.

Limited information content in scalar measurements and spurious signals from temporal variations make it difficult to estimate parameters beyond surface map localization of small-

scale anomalies for small-scale surveys. Depth estimates for large-scale anomalies can be made from total field data by relying on graphical and analytical methods developed by Peters [115], Smith [116], and Bott and Smith [117]. Peter’s Method is suboptimal for small-scale surveys involving small-scale anomalies; however, since it requires many simplifying assumptions to be made about the target, including having uniform magnetization and vertical sides with a large depth extent, and requires magnetic contour maps whose spacings are equal and proportional to the depth of the source [118]. Maximum depth rules derived by Bott and Smith are best suited when the anomaly is caused by an isolated, single source, which is not the case for many minefields. Euler’s Equation and Werner Deconvolution can approximate depths when there are many sources (or volumes of sources) of simple shape, but these methods are best suited for large surveys with long profiles [119, 120]. Additionally, Euler’s method requires information about the anomaly directional gradients, which are not directly sampled using scalar magnetometers but instead calculated in the Fourier domain using multiple filtering operations that can be unstable depending on the orientation of the anomaly magnetization vector and ambient field vector [118].

Scalar field measurements, by nature, contain less information about anomaly magnetic susceptibility, shape, and orientation, making interpretations difficult as their inversions are highly non-unique when used in isolation. Some of the non-uniqueness can be reduced by jointly inverting with other data modalities or introducing known priors [121–123]. For small-scale, small-anomaly targets such as landmines and UXO, there is an additional challenge of distinguishing the targets of interest from metallic clutter or shrapnel, and inversion methods using total field data rely on approximating the octupole component of a simplified spheroid, or assuming an arbitrarily axially symmetric body and attempting to recover the dipole and quadrupole moments from it [124, 125]. While transformations can be applied to total field measurements to ap-

proximate vector components, the operations become unstable depending on the direction of the ambient field [118]. Additionally, any instrumentation noise will map to spurious artifacts during scalar-to-vector component transformations [126].

This motivates the use of vector magnetometers, which can provide up to three times as much information about buried objects, including sensitivity to their orientation and shape, more directly than using total field data transformations to approximate vector components.

3.1.2.2 Vector Magnetometry

Vector magnetometers provide greater utility, are lightweight and field-rugged, and provide more information than scalar magnetometers by measuring vector components directly. Like scalar magnetometers, vector magnetometers suffer from the same temporal-to-spatial variation mapping and sensor-height-to-signal-amplitude limitations as scalar magnetometers. Additionally, they are exceedingly sensitive to heading errors (time-varying rotations of the magnetometer with respect to the desired pointing orientation) and have limitations due to fixed non-orthogonality and offset errors [43, 127]. Field calibration efforts have been made to account for some of the sensitivities and offset errors [127–130]. Still, sensor heading (pointing) errors remain difficult to compensate for in human-held and UAV-based magnetometer systems [131].

Figure 3.3 shows signal error as the percentage of expected magnetic anomaly for the same AT mine in Figure 3.2, measured by a triaxial fluxgate magnetometer from 0.5 *m* height, as a function of the pointing (heading) error of the sensor. B_x , B_y , and B_z represent the vector components of the expected magnetic anomaly, and ambient magnetic inclination and declination values are for College Park, MD. Even very small pointing errors ($< 1^\circ$) result in spurious mag-




Figure 3.3: Signal error as the percentage of expected magnetic anomaly for a metal AT mine buried at 15 *cm* depth, measured with a triaxial fluxgate magnetometer from 0.5 *m* height, as a function of the pointing (heading) error of the sensor. $B_x = 150nT$, $B_y = 200nT$, $B_z = 400nT$, declination = -10.5° , inclination = 67.12° (matching College Park, MD). Very small pointing errors produce spurious magnetic signals large enough to obscure those from an AT mine, increasing the chances of false flags.

netic signals that are large enough in amplitude to mimic those from the AT mine, increasing the chances of false flags. This relationship extends to any small, dipole-like target, making precise pointing crucial for mapping their magnetic signatures.

However, maintaining sub-degree pointing accuracy is extremely difficult for human-held fluxgate magnetometers and is practically untenable for aerial systems, including UAVs. While the use of MEMS accelerometers can help determine the orientation of fluxgate magnetometer sensor heads for making corrections, achieving accuracies $< 1^\circ$ is difficult in a dynamic environment [132], and vibrations aboard crafts such as UAVs produce additional inaccuracies that can create spurious signals at the same scale as targets of interest.

3.1.2.3 Magnetic Gradiometry

Magnetic gradiometry largely solves the problem of temporal-to-spatial mapping since gradiometers are sensitive to nearby variations and relatively insensitive to distant variations [133]. This is because the correlated signals produced by cultural and space weather can be effectively suppressed by differencing signals recorded at nearby magnetometers. For an infinitesimal triaxial gradiometer, this results in a nine-component tensor (IMGT), $\hat{G}$, represented as

$$\hat{G} = \vec{\nabla} \vec{B} = \partial_k B_i \quad (3.1)$$

Since curl and divergence are both zero outside of source regions, according to Maxwell's equations, there will only be five independent elements and the IMGT can then be written as a symmetric matrix

$$\hat{\mathbf{G}} = \begin{pmatrix} B_{xx} & B_{xy} & B_{xz} \\ B_{xy} & B_{yy} & B_{yz} \\ B_{xz} & B_{yz} & -B_{xx} - B_{yy} \end{pmatrix} = \hat{\mathbf{G}}^T \quad (3.2)$$

While magnetic gradiometry eliminates distant noise sources and removes the need for secondary base stations to correct survey data, it remains sensitive to heading errors. This motivated the introduction of the magnetic gradient tensor invariants that are insensitive to rotation [134–136]. Most applications of magnetic gradiometry assume the idealized IMGT which will be, by definition, symmetric and traceless [118]. In reality, physical magnetic gradiometer systems rely on finite-difference approximations of the IMGT, and, as a result, there will be errors between the theoretical and finite-differences magnetic gradient tensor (FDMGT) [85]. In practice, gradients are approximated by the separation-normalized difference between two or more magnetometers. Using the methods of Xu et al. [85] for a regular tetrahedron configuration (shown in Figure 3.4(a)), the FDMGT, $\mathbf{G}$, will then be of the form

$$\mathbf{G} = \frac{1}{d\sqrt{6}} \begin{pmatrix} \sqrt{2}(2B_{1x} - B_{2x} - B_{3x}) & \sqrt{2}(2B_{1y} - B_{2y} - B_{3y}) & \sqrt{2}(2B_{1z} - B_{2z} - B_{3z}) \\ \sqrt{2}(B_{1x} + B_{2x} - 2B_{3x}) & \sqrt{2}(B_{1y} + B_{2y} - 2B_{3y}) & \sqrt{2}(B_{1z} + B_{2z} - 2B_{3z}) \\ 3B_{4x} - B_{1x} - B_{2x} - B_{3x} & 3B_{4y} - B_{1y} - B_{2y} - B_{3y} & 3B_{4z} - B_{1z} - B_{2z} - B_{3z} \end{pmatrix} \quad (3.3)$$

where B_{1x} corresponds to the x-component of magnetometer 1, B_{1y} corresponds to the y-component of magnetometer 1, and so on.




Figure 3.4: (a) Schematic of a regular tetrahedron-based triaxial magnetic gradiometer. 1, 2, 3, and 4 are the positions of triaxial fluxgate magnetometers, each at the vertices of a regular tetrahedron. The side length is labeled as d . Each magnetometer will have x , y , and z components (B_{1x} , B_{1y} , B_{1z} , B_{2x} , etc). (b) TetraMag magnetic gradiometer system. MEDA FVM400 triaxial fluxgate magnetometer sensor heads are in 3D printed enclosures at the four vertices of a regular tetrahedron with $d = 0.5 \text{ m}$ side length.

3.1.2.4 Magnetic Field Finite Differences

The differences between the FDMGT and IMGMT capture information about magnetic field lines bending as they approach their dipolar source and can provide useful information about measured anomalies. Here we will explore how these asymmetries arise when using finite differences for a vertically separated gradiometer to first- and second-order in the dimensionless ratio of magnetometer size to target distance, ϵ , and how these errors relate to the distance from the dipole. The magnetic field $\vec{B}$ at location $\vec{r}$ away from source regions can be written as the gradient of a scalar potential, Φ

$$\vec{B}(\vec{r}) = \vec{\nabla}\Phi(\vec{r}). \quad (3.4)$$

For a point dipole with magnetic moment $\vec{m}$, the magnetic field becomes, in SI units:

$$\vec{B} = \vec{\nabla} \left(\vec{m} \cdot \vec{\nabla} C_m r^{-1} \right) \quad (3.5)$$

where $C_m = \mu_0/4\pi$, μ_0 is the magnetic permeability of free space, and $r = |\vec{r}| = \sqrt{\sum_k x_k^2}$.

Remembering useful results that $\partial_i r = x_i/r$ and $\partial_i r^{-1} = -x_i r^{-3}$, and considering the special case of a purely vertical magnetic moment, $\vec{m} = \sum_k m_0 \delta_{k3} \hat{e}_k$, the expression for the magnetic field simplifies to:

$$\vec{B}(\vec{r}) = -m_0 C_m \vec{\nabla} (x_3 r^{-3}) = \frac{3m_0 C_m x_3}{r^5} \sum_i \left(x_i - \frac{r^2}{3x_i} \delta_{i3} \right) \hat{e}_i \quad (3.6)$$

The magnetic gradient tensor $\vec{\nabla} \vec{B} = \partial_k B_i = \partial_k \partial_i \Phi$ is inherently symmetric because of the symmetry of second derivatives. However, when we measure the magnetic gradients, we actually measure the finite differences in the components of the magnetic field, B_i , across some physical separations $\vec{l}$. We can write this magnetic finite differences matrix, $\mathbf{G}$ as:

$$G_{ij} = \frac{1}{|\vec{l}|} \left[B_i \left(\sum_k x_k + l_j \delta_{kj} \hat{e}_k \right) - B_i \left(\sum_k x_k \hat{e}_k \right) \right] \quad (3.7)$$

Let us consider taking the finite difference across a vertical separation, i.e. $\vec{l} = l_3 \hat{e}_3$, of the $\hat{x}_1$ component of the magnetic field, and write the physical size of the magnetic gradiometer in terms of the vertical distance from the source, i.e. $l_3 = \epsilon x_3 = |\vec{l}|$.

$$G_{13} = \frac{3m_0 C_m}{\epsilon x_3} \left[\frac{1 + \epsilon}{(\sum_k x_k^2 (1 + \epsilon \delta_{k3})^2)^{5/2}} - \frac{1}{(\sum_k x_k^2)^{5/2}} \right] x_3 x_1 \quad (3.8)$$

When the scale of the magnetic gradiometer is small compared to the distance from the magnetic anomaly being characterized, which is the case in most applications of magnetic gradiometry, $\epsilon \ll 1$, we need only account for terms of the first order in ϵ . However, in our use case of characterizing magnetic anomalies that are small in amplitude and small in spatial extent, ϵ can approach unity, and we cannot neglect $\mathcal{O}(\epsilon^2)$ terms. Therefore, we proceed by expanding the terms to second order in ϵ using the binomial approximation: $(1 + \epsilon)^\alpha \approx 1 + \alpha\epsilon + (\alpha/2)(\alpha - 1)\epsilon^2$. After some algebra, the first term in equation 3.8 becomes:

$$\begin{aligned} & \frac{(1 + \epsilon)}{r^5} [1 + (2\epsilon + \epsilon^2)x_3^2/r^2]^{-5/2} \approx \\ & \frac{1}{r^5} \left[1 + \left(1 - 5\frac{x_3^2}{r^2}\right)\epsilon - \frac{5}{2}\frac{x_3^2}{r^2} \left(3 - 7\frac{x_3^2}{r^2}\right)\epsilon^2 + \mathcal{O}(\epsilon^3) \right] \end{aligned} \quad (3.9)$$

The entire expression for the vertical finite difference of the $\hat{x}_1$ component of the magnetic field (3.8) then simplifies to:

$$G_{13} = \frac{3m_0C_m}{r^5} \left[\left(1 - 5\frac{x_3^2}{r^2}\right) - \frac{5}{2}\frac{x_3^2}{r^2} \left(3 - 7\frac{x_3^2}{r^2}\right)\epsilon + \mathcal{O}(\epsilon^2) \right] x_1 \quad (3.10)$$

Now, we turn our attention to the finite difference in the $\hat{x}_i$ direction (i.e. $\vec{l} = l_i\hat{e}_i = \epsilon x_i\hat{e}_1$) of the vertical component of the magnetic field:

$$G_{3i} = \frac{3m_0C_m}{\epsilon x_i} \left[\frac{x_3 - \frac{1}{3x_3} \sum_k x_k^2 (1 + \epsilon\delta_{ki})^2}{(\sum_k x_k^2 (1 + \epsilon\delta_{ki})^2)^{5/2}} - \frac{x_3 - \frac{r^2}{3x_3}}{(\sum_k x_k^2)^{5/2}} \right] x_3 \quad (3.11)$$

Simplifying the summation terms using r

$$\sum_k x_k^2 (1 + \epsilon\delta_{ki})^2 = r^2 \left[1 + (2\epsilon + \epsilon^2)\frac{x_i^2}{r^2} \right] \quad (3.12)$$

and using the binomial approximation to second order in ϵ

$$\left(\sum_k x_k^2 (1 + \epsilon \delta_{ki})^2 \right)^{-5/2} \approx r^{-5} \left[1 - 5 \frac{x_i^2}{r^2} \epsilon - \frac{5x_i^2}{2r^2} \left(1 - 7 \frac{x_i^2}{r^2} \right) \epsilon^2 + \mathcal{O}(\epsilon^3) \right] \quad (3.13)$$

we can approximate equation 3.11 to read:

$$G_{3i} \approx \frac{3m_0 C_m}{\epsilon x_i r^5} \left[\left(x_3 - \frac{r^2}{3x_3} - (2\epsilon + \epsilon^2) \frac{x_i^2}{3x_3} \right) * \left[1 - 5 \frac{x_i^2}{r^2} \epsilon - \frac{5x_i^2}{2r^2} \left(1 - 7 \frac{x_i^2}{r^2} \right) \epsilon^2 \right] - \left(x_3 - \frac{r^2}{3x_3} \right) \right] x_3 \quad (3.14)$$

After some algebra, the expression above simplifies further:

$$G_{3i} \approx \frac{3m_0 C_m}{\epsilon x_i r^5} \left[\left(\frac{r^2}{3x_3} - x_3 \right) \left[5 \frac{x_i^2}{r^2} \epsilon + \frac{5x_i^2}{2r^2} \left(1 - 7 \frac{x_i^2}{r^2} \right) \epsilon^2 \right] - (2\epsilon + \epsilon^2) \frac{x_i^2}{3x_3} \left(1 - 5 \frac{x_i^2}{r^2} \epsilon \right) \right] x_3 \quad (3.15)$$

Rearranging by order of ϵ , we arrive at:

$$G_{3i} = \frac{3m_0 C_m}{r^5} \left[\left(1 - 5 \frac{x_3^2}{r^2} \right) + \left(\frac{1}{2} - \frac{15x_i^2}{6r^2} - \frac{5x_3^2}{2r^2} + \frac{35x_i^2 x_3^2}{2r^4} \right) \epsilon + \mathcal{O}(\epsilon^2) \right] x_i \quad (3.16)$$

Comparing equations 3.10 and 3.16 we can see that when $\epsilon \ll 1$, the two are equal, as we expected from the symmetry of second derivatives of the scalar potential. However, when characterizing nearby, small-scale anomalies, $\epsilon \approx 1$ and the higher order terms produce differences between the two equations. This manifests as an asymmetry in the finite different magnetic gradient matrix. The amplitude of the asymmetry is proportional to the amplitude of the higher-order spatial derivatives of $\vec{B}$. The ratio of the second to first-order spatial derivatives of a vector field is related to the curvature of the field, with larger curvatures corresponding to stronger second-order




Figure 3.5: The $\hat{x}_3$ finite difference of the $\hat{x}_1$ component of the magnetic field (G_{13} , black solid line) and the $\hat{x}_1$ finite difference of the $\hat{x}_3$ component of the magnetic field (G_{31} , red solid line) are poorly approximated when only the first term in equations 3.10 (black dashed), 3.16 (red dashed) is used. However, the inclusion of second-order terms in ϵ (dot-dashed lines) improves accuracy. Note that the dashed lines are parallel because of the symmetry of $\vec{\nabla} \vec{B}$

spatial derivatives. When we express the $FDMGT$ as a sum of a symmetric and antisymmetric matrix, the relative size of the antisymmetric matrix quantifies the local curvature of $\vec{B}$. Since higher-order spatial derivatives – i.e. terms higher order in ϵ – fall off more quickly with distance r , the curvature of $\vec{B}$ and the size of the antisymmetric component of $FDMGT$ can serve as a proxy for the gradiometer’s proximity to a localized magnetic anomaly.

3.2 Asymmetry of the Finite-Difference Magnetic Gradient Tensor

As we found in the previous section, differences between the IMGT, $\hat{\mathbf{G}}$, and FDMGT, $\mathbf{G}$, will appear as deviations from the expected symmetry of the IMGT. To highlight these deviations, it is useful to represent FDMGT as a sum of its symmetric part, $\mathbf{G}_\Gamma 30803$, and its antisymmetric part, $\mathbf{G}_\Gamma 30785$,

$$\mathbf{G} = \mathbf{G}_S + \mathbf{G}_\Gamma 30785 = \frac{\mathbf{G} + \mathbf{G}^T}{2} + \frac{\mathbf{G} - \mathbf{G}^T}{2} \quad (3.17)$$

Functionally, the antisymmetric part represents errors that arise from the physical size and geometry of the magnetic gradiometer.

The Frobenius norm provides a useful way of comparing matrices, being a measure of the average response along mutually orthogonal directions in space. It is analogous to an L2 or Euclidian norm but for a matrix with a rank greater than 1 [137]. The Frobenius norm of the symmetric part of the FDMGT, $\|$

$\mathbf{G}_S\|_F$, and the antisymmetric part, $\|\mathbf{G}_\Gamma 30785\|_F$ can be expressed as [137]

$$\|\mathbf{G}_{A,S}\|_F = \sqrt{\sum_{i=1}^3 \sum_{j=1}^3 |g_{A,S,ij}|^2} \quad (3.18)$$

where $g_{S,ij}$ and $g_{A,ij}$ are the elements of $\mathbf{G}_\Gamma 30803$ and $\mathbf{G}_\Gamma 30785$, respectively.

We use the Frobenius norm of the symmetric and antisymmetric components of FDMGT to quantify how the finiteness (side length) of the gradiometer affects deviations from the symmetry of IMGT. In Figure 3.6, we plot $\|\mathbf{G}_\Gamma 30803\|_F$ and $\|\mathbf{G}_\Gamma 30785\|_F$ for a vertically-oriented dipole measured from a height of 1.5 m. For small side lengths, the Frobenius norm is dominated by the




Figure 3.6: Symmetric (red) and antisymmetric (blue) components of FDMGT exhibit different behavior with increasing gradiometer side lengths. $\|\mathbf{G}_{\Gamma 30785}\|_F$ grows with the square of the side length (dotted line), while $\|\mathbf{G}_{\Gamma 30803}\|_F$ remains unchanged so that their ratio also grows with the square of the side length. For small side lengths, the antisymmetric component is negligible but approaches the symmetric component for realistic gradiometer sizes.

symmetric part of the tensor. As the side length increases to realistic values expected for physical gradiometer systems, the antisymmetric part grows and is eventually the same amplitude as the symmetric part. As a result, the ratio of $\|\mathbf{G}_{\Gamma 30785}\|_F / \|\mathbf{G}_{\Gamma 30803}\|_F$ grows exponentially with the square of the length/distance ratio, which may be leveraged for determining the depth of a detected anomaly.

3.3 Invariants of the Magnetic Gradient Tensor

To minimize the impact of rotations and heading errors for a finite triaxial gradiometer and explore which properties of the FDMGT may be useful for small-scale anomaly detection and source parameter characterization, we seek to utilize invariants of the FDMGT. Invariants of interest for this study include previously described Frobenius norms and eigenvalues, determinants,

and normalized source strengths (NSS) of the FDMGT. Eigenvalues are scalars used to transform eigenvectors, which distill key information in a tensor while maintaining the relationships between individual components. The FDMGT can be written in terms of its eigenvalues [138]

$$\mathbf{G} = \mathbf{V}^T \begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{pmatrix} \mathbf{V} \quad (3.19)$$

where matrix $\mathbf{V}$ contains mutually orthogonal eigenvectors $[\vec{v}_1 \ \vec{v}_2 \ \vec{v}_3]$ and $\lambda_1, \lambda_2, \lambda_3$ are the corresponding eigenvalues ($\lambda_1 > \lambda_2 > \lambda_3$).

If we consider the magnetic gradient tensor to define a linear transformation, its determinant is the signed scaling factor for its volume and is invariant to rotation. However, this invariance is not guaranteed for the FDMGT. To explore how the determinant will change with errors introduced by the physical, regular tetrahedral magnetic gradiometer system, we need only consider the determinant of the symmetric part of the FDMGT, $|\mathbf{G}_{\Gamma 30803}|$, which can be expressed as

$$|\mathbf{G}_{\Gamma 30803}| = \lambda_1 \lambda_2 \lambda_3, \quad (3.20)$$

because the determinant of the antisymmetric part, $|\mathbf{G}_{\Gamma 30785}|$, is zero by construction.

In addition to the traditional invariants described above, we can combine eigenvalues in various ways to explore their sensitivity to target parameters. For this study, we focus on two eigenvalue combinations, λ_{C1} and λ_{C2} , which can be expressed as

$$\lambda_{C1} = \lambda_1^2 + \lambda_2^2 + \lambda_3^2 \quad (3.21)$$

$$\lambda_{C2} = \lambda_1^3 + \lambda_2^3 + \lambda_3^3 \quad (3.22)$$

The NSS is a rotational invariant of the FDMGT and can be expressed as a combination of eigenvalues λ_1 , λ_2 , and λ_3 [139]

$$NSS = \sqrt{-\lambda_2^2 - \lambda_1\lambda_3} \quad (3.23)$$

The NSS peaks over compact 3D sources [139–141], which is particularly interesting to this study for locating and characterizing small-scale, small-amplitude targets such as landmines and UXO.

3.4 Data

Here, we present the data acquisition system, synthetic data and field data.

3.4.1 Magnetic Gradiometer System (TetraMag)

We have developed a magnetic gradiometer called TetraMag, using four MEDA FVM400 triaxial fluxgate magnetometers whose data directly sample the full FDMGT. TetraMag is designed to detect small-scale, small-amplitude targets from very low altitudes (~ 1 m AGL) in the shallow critical zone. Vector magnetometers were chosen as the base sensors since gradient magnetometers constructed from scalar sensors suffer from significant limitations described in the vector magnetometry section. Their suboptimal performance is particularly important when used at low altitudes, as the vertical gradient cannot be calculated reliably due to difficulties

controlling altitude and the proximity to metallic objects [142, 143].

The four triaxial magnetometer sensor heads are in 3D-printed housings at the vertices of a regular tetrahedron, shown in Figure 3.4(b). A regular tetrahedron was chosen based on the work of Xu et al. [85], who showed that a regular tetrahedron configuration outperforms a right-angled tetrahedron for triaxial gradiometry purposes. Since our expected targets (landmines and UXO) are of spatial scales on the order of $0.05 - 0.2 \text{ m}$, we selected the side length of $d = 0.5 \text{ m}$.

Carbon fiber tubing connects each sensor head to a 3D-printed central hub. Aluminum supports help to control structural modes and allow TetraMag to be rigidly mounted to the survey platform of choice. Figure 3.4(b) shows TetraMag mounted on a rolling aluminum scaffold for test surveys.

3.4.2 Synthetic Data

The forward model to generate synthetic data uses the physical geometry of TetraMag to calculate the expected vector components at each vertex location and then calculate the FDMGT, G , for a chosen sensor measuring height. The magnetic vector components for the dipole source were calculated using the methods presented by Smellie [144] and modified by Oruç [83]. We assume magnetization is due to Earth's field only so that the components can be expressed as [83].

$$B_x = 2p_0 \frac{r^2 l - 2r_x(r_x l + r_y m - z_0 n)}{r^4} \quad (3.24)$$

$$B_y = 2p_0 \frac{r^2 m - 2r_y(r_x l + r_y m - z_0 n)}{r^4} \quad (3.25)$$

$$B_z = -2p_0 \frac{r^2 n - 2z_0(r_x l + r_y m - z_0 n)}{r^4} \quad (3.26)$$

where p_0 is the magnetic moment, z_0 is the depth below the plane of observation, l , m , and n are the direction cosines that depend on the magnetic inclination (I) and declination (D), and r is the distance between the dipole and measurement location, with r_x and r_y components in the x and y directions, respectively.

At each measurement location, we project the measured components onto the centroid of the tetrahedron and calculate the FDMGT, $\mathbf{G}$, (Equation 3.3). From the full FDMGT, which we visualize as pixels centered on the gradiometer centroid, we calculate the values of each invariant. Figure 3.7 shows a schematic visualization for a dipole target buried at 15 *cm* depth and oriented vertically ($I = 90^\circ$, $D = 0^\circ$), measured from 0.5 *m* height. Panel (a) shows the magnetic field lines and TetraMag to scale, (b) shows the survey path in map view and centroid position where invariants are calculated, and (c) shows the visualization of a calculated example invariant at each pixel location. The pixels are left in gridded format for demonstration purposes only.

3.4.3 Field Data

Validation data were collected using TetraMag mounted on an aluminum scaffold over a bar magnetic (meant to approximate a metallic AP landmine) buried at 15 *cm* depth in a test sand bed and measured from 55 *cm* height. The test bed consists of a 300-gallon polyethylene containment tray filled with sand, located at the Paint Branch Turfgrass Research Facility in College Park, Maryland. The aluminum scaffold has MRI-safe wheels and an aluminum 8020 slider, allowing TetraMag to move in two directions, resulting in a lawnmower-style survey pattern shown in




Figure 3.7: Survey geometry and example schematic. (a) Magnetic field lines for a vertical dipole buried at 15 *cm* depth and TetraMag, to scale (lower left), with orientations of B_x , B_y , and B_z for magnetometer 1. (b) Lawnmower-style survey patterns followed by TetraMag, with gradients, G_{T30790} , referenced to its centroid (black star). (c) An example visualization for one of the G_{T30790} invariants, with color of each pixel representing a chosen G_{T30790} invariant at that location.

Figure 3.7 (b). Two 1-D LiDAR units were used to keep track of the position of TetraMag during a survey and tied to individual magnetometer measurements via a shared clock on a Raspberry Pi 4 that simultaneously collected TetraMag and position data. The data were interpolated onto a spatial grid, and the components of the FDMGT were calculated at each pixel point using the methods described above for a regular tetrahedral gradiometer and presented in Xu et al. [85].

3.5 Results

In this section, we present synthetic data for a vertically oriented dipole, followed by synthetic data that reflect the survey area's magnetic inclination, declination, estimated target magnetic moment, and TetraMag tilt. We then validate the expected results with the field data described above.

We can visualize the synthetic data in the way shown in Figure 3.7 to explore how the IMGT and FDMGT will differ from one another as we rotate TetraMag. Figure 3.8, columns 1-3, shows how the patterns of λ_1 , λ_2 , and λ_3 vary for IMGT and FDMGT when TetraMag is oriented parallel to the ground (IMGT No rot., FDMT No rot.) and when it is rotated 30° with respect to the x-axis (FDMGT 30° rot.), each measured from 0.5 m height. The dipole target is indicated by a white asterisk and buried at 15 cm depth with vertical orientation ($I = 90^\circ$, $D = 0^\circ$). As expected, the visualized patterns for IMGT do not change with 30° rotation since the measurements are close to infinitesimal (side length = 0.000001 m), so their plots are omitted. The target is localized with an anomalous high lobe for λ_1 and λ_2 , and an anomalous low lobe for λ_3 . For the FDMGT case, however, the target is localized by an anomalous low lobe for λ_1 , λ_2 , and λ_3 . λ_1 is surrounded by two positive high lobes which change orientation as TetraMag is

tilted 30° . It should be noted that the sign of the eigenvalues is arbitrary, but that its magnitude differs significantly from the surrounding survey area.

Figure 3.8, columns 4-5, shows how determinant values for the full MGT, as well as its symmetric and antisymmetric parts, vary for the same TetraMag orientations. An anomalous high lobe for IMGT localizes the target, and the antisymmetric part, $|\mathbf{G}_T30785|$, is zero since the infinitesimal case is traceless and symmetric. For the FDMGT case, the target is localized by an anomalous low lobe whose $|\mathbf{G}_T30803|$ part reveals surrounding anomalous high lobes. The locality of the anomalous high lobe changes as TetraMag is tilted, and $|\mathbf{G}_T30785|$ is non-zero since the FDMGT is an approximation based on a side length of 0.5 m .

Figure 3.8, columns 6-8, shows how Frobenius norm values for the full MGT, as well as its symmetric and antisymmetric parts, vary for the same set of TetraMag orientations. The target is localized by an anomalous high lobe for the IMGT and FDMGT cases, though the pattern and amount of visual diffusivity depend on TetraMag rotation for FDMGT. The $\|\mathbf{G}_T30785\|_F$ values for FDMGT arise from the 0.5 m TetraMag side length.

Figure 3.8, columns 9-11, shows how λ_{C1} , λ_{C2} , and NSS values vary for IMGT and FDMGT for the same TetraMag orientations. The target is localized by an anomalous high lobe for all three IMGT invariants, though the amplitude of λ_{C2} is significantly smaller than those from λ_{C1} and NSS . The target is localized by anomalous high lobes for λ_{C1} and NSS , and by a low lobe for λ_{C2} in the FDMGT case. When TetraMag is rotated in the FDMGT case, the patterns rotate slightly. The diameter of the high λ_{C2} lobe is proportional to the burial depth for both the IMGT and FDMGT cases.

Figure 3.9 shows the average expected (Synth.) and measured (Field) values of B_x , B_y , and B_z . Streak-like, spurious features in the field data are caused by inconsistent survey spac-




Figure 3.8: Maps of the eigenvalues (columns 1-3) and invariants (columns 4-11) calculated from synthetic data at two TetraMag orientations for the full IMGT and FDMGT, $G_{\Gamma 30822}$, as well as the latter's symmetric, $G_{\Gamma 30803}$, and antisymmetric, $G_{\Gamma 30785}$, parts. Eigenvalues λ_1 , λ_2 , and λ_3 are shown in columns 1-3. Determinant variations, $|G_{\Gamma 30822}|$ and $|G_{\Gamma 30803}|$, are shown in columns 4-5. Frobenius norm variations, $\|G_{\Gamma 30822}\|_F$, $\|G_{\Gamma 30803}\|_F$, and $\|G_{\Gamma 30785}\|_F$, are shown in columns 6-8. Variations of eigenvalue combinations 1 (λ_{C1}) and 2 (λ_{C2}) and NSS are shown in columns 9-11.

ing, which maps into spatial features during interpolation (varying between 5-11 *cm* in the *y*-direction). Additionally, magnetic noise from nearby ferromagnetic clutter has been piecewise removed from the B_y data, ghosting from which can be seen near $x = 50$ *cm*.

Figure 3.10, columns 1-3, shows eigenvalues λ_1 , λ_2 , and λ_3 for field and synthetic data. There is a slight horizontal shift in the anomalous positive lobe of λ_1 field data compared to the expected values in the synthetic data. However, the target still appears along the edge of the lobe. The target is well-localized at the center of the anomalous positive lobe of λ_2 for both the field and synthetic data. An additional anomalous negative lobe is seen in the field data that is not present in the synthetic data. Target localization is consistently in the center of the anomalous negative lobe in λ_3 for both synthetic and field data, though the pattern is slightly more diffuse for the field data.

Figure 3.10, columns 4-5, shows the determinant of the full FDMGT ($|G_{\Gamma 30822}|$), as well as its symmetric part ($|G_{\Gamma 30803}|$). There is good agreement between the expected patterns for $|G_{\Gamma 30822}|$ in the synthetic and field data, with the target being localized in the center of the anomalous negative lobe in $|G_{\Gamma 30803}|$. There is an additional positive anomaly seen in the field data that is not present in the synthetic data.

Figure 3.10, columns 6-8, shows the Frobenius norm of the full FDMGT, $\|G\|_F$, as well as its symmetric ($\|G_{\Gamma 30803}\|_F$) and antisymmetric ($\|G_{\Gamma 30785}\|_F$) parts. The expected patterns from the synthetic data match the field data well, especially for $\|G\|_F$ and $\|G_{\Gamma 30803}\|_F$. The target is localized in the center of the anomalous positive lobe for both $\|G\|_F$ and $\|G_{\Gamma 30803}\|_F$, with finer resolution seen in $\|G_{\Gamma 30803}\|_F$ for both synthetic and field data, though the lobe appears more elongated in the field data which is revealed to be a secondary positive lobe upon inspection in $\|G_{\Gamma 30803}\|_F$. The target is localized at the edge of a positive anomaly lobe in




Figure 3.10: Maps of eigenvalues (columns 1-3) and invariants (columns 4-11) for field FDMGT data (Field) collected by TetraMag and synthetic FDMGT data (Synth.) calculated using TetraMag's geometry, tilt, and local magnetic conditions. Eigenvalues λ_1 , λ_2 , and λ_3 are shown in columns 1-3. Determinant variations, $|G_f|$ and $|G_s|$, are shown in columns 4-5. Frobenius norm variations, $\|G_{T-30822}\|_F$, $\|G_{T-30803}\|_F$, and $\|G_{T-30785}\|_F$, are shown in columns 6-8. Eigenvalue combinations 1 (λ_{C1}) and 2 (λ_{C2}) and NSS variations are shown in columns 9-11.

$\|\mathbf{G}_{\Gamma 30785}\|_F$ for both synthetic and field data, though two additional positive anomaly lobes are seen in the field data that are not present in the synthetic data.




Figure 3.9: Average B_x , B_y , and B_z values for synthetic data and corresponding field data measured by TetraMag at 55 cm height. The dipole target is indicated by a black asterisk and buried 15 cm in a test sand bed in College Park, MD. Synthetic data match the field magnetic inclination and declination values. Noise from magnetic clutter on the test scaffold has been piecewise removed from the field B_y values, ghosting from which can be seen in the plot.

Figure 3.10, columns 9-11, shows eigenvalue combinations λ_{C1} , λ_{C2} , and normalized source strength NSS for field and data. The target is localized in the center of a positive anomaly lobe for λ_{C1} and a negative anomaly lobe for λ_{C2} in both synthetic and field data, though λ_{C1} is more elongated than the expected circular lobe in synthetic data. The positive anomaly lobe seen to the northwest of the negative lobe in λ_{C2} is present in both synthetic and field data. There is a slight horizontal offset between the positive anomaly lobe in the synthetic data and field data, and an additional positive lobe is seen in the field data that is unseen in the synthetic data.

3.6 Discussion

3.6.0.1 The Impact of Finiteness on FGMGT Invariants

The finiteness of TetraMag introduces an antisymmetric component to the FDMGT.

As discussed in the magnetic gradiometry section, $\|\mathbf{G}_{\Gamma 30785}\|_F$ grows with increasing side length and decreasing distance from a small-scale magnetic anomaly. Its ratio to $\|\mathbf{G}_{\Gamma 30803}\|_F$ grows with the square of the side length, as shown in Figure 3.6. However, the dependence of the ratio on location exhibits more complicated patterns, reflecting the magnetization of the target. In Figure 3.11, we display $\|\mathbf{G}_{\Gamma 30785}\|_F / \|\mathbf{G}_{\Gamma 30803}\|_F$ at two flying heights, 0.5 *m* (top row), corresponding to a ground-based survey, and 1.5 *m* (bottom row), corresponding to a low-altitude drone-based survey; side length is set equal to the flying height in both cases. The columns show how these patterns change with differing magnetization values. When the magnetization is oriented vertically (right column), for example, the resulting pattern localizes the anomaly with a small radius ring of high-amplitude values encircling the target location. The radius of this ring remains unchanged with changing flying height, suggesting it may be leveraged to characterize its source parameters. This symmetry does not persist for horizontally oriented or dipping magnetization vectors, as seen in the center and left columns, which show lobate patterns.

The side length of TetraMag also plays a significant role in the expected invariant patterns derived from the FDMGT, which can be seen from the synthetic data and validated by the field

data. As the side length of the regular tetrahedral gradiometer increases, new features such as additional positive or negative anomaly lobes can be seen surrounding the dipole target area (present in λ_1 , $|\mathbf{G}_\Gamma 30803|$, $\|\mathbf{G}_\Gamma 30803\|_F$, and λ_{C2}). These structures reveal information that can help estimate source parameters such as magnetic moment and burial depth in addition to map localization. These features arise from the finiteness of the FDMGT producing an antisymmetric part of the gradient tensor.

While the synthetic IMGT data presented here are rotationally invariant, consistent with the literature [83, 131, 138, 140], the FDMGT is still subject to rotational variations of the physical gradiometer which produce small changes in the expected pixel-to-pixel invariant pattern. The invariant pattern differences when TetraMag is rotated are most apparent in λ_1 , λ_3 , $|\mathbf{G}_\Gamma 30803|$, $\|\mathbf{G}\|_F$, $\|\mathbf{G}_\Gamma 30803\|_F$, $\|\mathbf{G}_\Gamma 30785\|_F$, and λ_{C2} .

For $|\mathbf{G}_\Gamma 30803|$, rotation of TetraMag results in a rotation of the vertices of the triangular anomaly pattern, but the location of the central negative lobe is unchanged. Similarly, the central negative lobe is unchanged for λ_1 and λ_{C2} , but there is an apparent shift in the surrounding positive anomaly lobe locations. The lateral extent of the central negative lobe is reduced with TetraMag rotation for λ_3 , but the target remains centered within 2 *cm* of the non-rotated pattern. In each of these cases, the target can still be localized well even with changes in surrounding lobe patterns.

Rotation of TetraMag results in a preferential amplitude increase for one corner of the positive anomaly lobe for $\|\mathbf{G}_\Gamma 30803\|_F$, $\|\mathbf{G}_\Gamma 30803\|_F$, and $\|\mathbf{G}_\Gamma 30785\|_F$. For $\|\mathbf{G}_\Gamma 30785\|_F$ in particular, this preferential brightening gives the appearance of a change in the overall invariant pattern, which could complicate interpretation if inspected in isolation, or in the event of large amplitude rotation of TetraMag.




Figure 3.11: Ratio of $\|\mathbf{G}_{\Gamma 30785}\|_F / \|\mathbf{G}_{\Gamma 30803}\|_F$ at two flying heights, 0.5 m (top row), corresponding to a ground-based survey, and 1.5 m (bottom row), corresponding to a low-altitude drone-based survey. The source is a magnetic dipole that is dipping (left), horizontal (center), or vertical (right). While a vertical magnetization produces the most symmetric $\|\mathbf{G}_{\Gamma 30785}\|_F / \|\mathbf{G}_{\Gamma 30803}\|_F$ pattern around the target location, dipping orientations produce more complicated, lobate patterns, but still peak above the target.

Invariants whose patterns are left largely unchanged when TetraMag is rotated include λ_{C1} , λ_2 , and NSS, emphasizing that not all MGT invariants are equally invariant to rotation. Separating the full FDMGT into its symmetric and antisymmetric parts, however, overcomes some of the errors introduced by having a physical, finite gradiometer and increases the localization resolution for the dipole target, especially for $|\mathbf{G}|$ and $\|\mathbf{G}\|_F$ but for the other invariants in this study as well. While each of the included FDMGT invariants localizes the target, λ_{C2} contains additional information about the target burial depth, demonstrating that some invariants are more useful depending on the desired source characterization.

3.6.0.2 Field Data Complications

Inconsistent survey y-direction spacing in the field (varying from 5 – 11 *cm*) mapped into spurious, streak-like features during the interpolation step of the FDMGT processing. This may have contributed to some of the errors between the expected synthetic invariant patterns and field data patterns. Additionally, the fiberglass scaffold rolling tracks used for the field survey were not level due to local topography, which resulted in a nearly 30° tilt of TetraMag with respect to the y-axis. All synthetic data have factored in this amount of tilt when calculating the FDMGT and its invariants.

3.7 Conclusions

This study illustrates how the finite size of magnetic gradiometers compared to the length-scales of magnetic anomalies due to small buried objects affects the symmetry and invariant properties of the finite-difference magnetic gradient tensor (FDMGT). This is in contrast to tra-

ditional assumptions of theoretical magnetic gradient tensors, which should be symmetric and whose invariants should not be affected by rotations. We showed that deviations from symmetry decrease with the square of distance from the magnetic target (represented as a dipole), but can exhibit complicated spatial patterns depending on the orientation of the magnetization of the target. Furthermore, we show how even when using the symmetric component of the FDMGT, rotational dependence of invariants persists.

We explored the utility and benefit of leveraging the symmetric and antisymmetric parts of the FDMGT for locating and characterizing small-scale, small-amplitude targets in the shallow critical zone. In particular, the Frobenius norm, NSS, and eigenvalue combinations λ_{C1} , λ_{C2} enable high-resolution localization of small magnetic targets and characterization of their source parameters such as burial depth and magnetization orientation. Building on theoretical frameworks from potential fields literature, we built and tested a regular tetrahedral triaxial magnetic gradiometer called TetraMag and demonstrated its ability to estimate the FDMGT, then compared each derived invariant's ability to localize the buried target.

The methodology presented here allows for rapid processing and interpretation of field data. Since most of the errors that arise between the IMGMT and FDMGT are wrapped up in the creation of an antisymmetric part of the MGT, utilizing the symmetric part obviates the need for some of the more rigorous calibration efforts traditionally applied to magnetic survey data.

This approach helps lay the foundation for UAV-based detection of small targets, and the relative stability of invariant patterns to rotations of TetraMag may enable the instrument to be rigidly mounted on the vehicle body. Accelerometer data could then be used to correct for TetraMag tilt during inversion calculations. This effectively shifts the correction task to the data analysis side rather than the field operational side, where gimbals and other equipment would

traditionally be used for the challenging task of always keeping the instrument level. Future investigations will include multiple dipole-like targets, as well as ferromagnetic clutter, to ensure the robustness of the approach. For machine learning (ML)-assisted inversion and source parameter estimation tasks, which UXO and landmine detection lend themselves to, synthetic training data should include rotations of the solid TetraMag structure to ensure all versions of the resulting invariant patterns are included as the network learns and updates.

Chapter 4: Machine Learning Methods for Detecting and Characterizing Landmines and UXO

In this chapter, we explore how machine learning (ML) methods may be leveraged to detect UXO in radargrams and to determine magnetic source parameters of targets using FDMGT invariants described in Chapter 3. We begin by introducing multi-layer perceptron and convolutional neural networks. We then detail an analog UXO site we set up and collected GPR data for, and discuss the design and optimization of a CNN for identifying whether a radargram contains UXO or not. We then extend the use of CNN to magnetic data by adding a regression layer to estimate source parameters such as burial depth and magnetic moment from detected targets. We systematically compare the utility of various types of magnetic data – total field, vector, gradient measurements, as well as various invariants of the gradient tensor. Finally, we discuss the utility and practicality of this technique in the context of traditional approaches, highlighting the advantages and limitations of including ML as an interpretation aid for near-surface geophysical data.

4.1 Classification and Regression Networks

Multi-Layer Perceptron (MLP) Multi-Layer Perceptron is a machine learning technique that was first introduced in Rosenblatt’s 1957 report, “The Perceptron, A Perceiving and Recognition




Figure 4.1: Architecture of a feed-forward multi-layer perceptron (MLP) network. (Source: [8])

Automaton,” which aimed to design an analog of the human brain for data analysis [145]. The network consisted of two fully connected layers: an input layer and an output layer which is a weighted sum of the input layer nodes with an applied Heaviside step activation function. The MLP builds on the Perceptron by including additional hidden layers and by introducing a nonlinear activation function with a nonzero gradient. Having a nonzero gradient is key for enabling the model to learn and update by optimizing an objective function via backpropagation. Backpropagation optimizes parameter values that control neuron activation by calculating the error function gradient with respect to network weights and updating the model as it trains [146]. Figure 4.2 shows a few example nonlinear activation functions used today [9]. The functions relevant for this study include the sigmoid function (Figure 4.2a), and the rectified linear unit (ReLU) function (Figure 4.2c). Sigmoid is a useful choice when the objective is to constrain the probability of an output; the output is then interpreted as a binary class label. ReLU is more appropriate for multi-class problems, when sparsity is desired and computational efficiency is a concern. The incorporation of stochastic gradient descent, a fast optimization technique that follows the nega-

tive gradient to minimize the error function, enables high-performance, time-efficient training of large datasets [147]. Gradient descent algorithms used in this study include root mean squared propagation (RMSProp), which increases the optimization speed by decreasing the number of function calculations to reach a local minimum, and Nadam, which is similar to RMSProp but with Nesterov momentum [148]. Nesterov momentum calculates the gradient at a position ahead in the direction of the accumulated momentum rather than at its current position, enabling faster convergence speeds. The efficiency and robustness of the produced models are largely influenced by one's choice of hyperparameters, which are configurable variables such as number of epochs (the number of times training data are passed through), learning rate (the step size at each iteration), dropout (number of nodes to ignore during training to avoid overfitting the model), batch size (the number of samples processed before updating the model), padding (preserving the spatial size of the input by adding zeros as necessary), and others.

MLPs can be used unaccompanied for classification problems, but are most often incorporated at the back end of other machine learning techniques in modern applications [149] and are included as the backend of the networks built in Sections 4.3 and 4.4. In Section 4.3, MLP is used as a binary classification mechanism; in 4.4, it provides the necessary data flattening and regression calculation capabilities to directly estimate source parameter values. Figure 4.1 shows an example of a feedforward MLP (Figure 1 from [8]). Here, four inputs enter the network and pass through two hidden layers, each with three neurons, before reaching an output layer, representing a binary classification problem.

Convolutional Neural Networks (CNN) Convolutional Neural Networks (CNNs) are prominent types of artificial neural networks, designed to replicate how we think humans process visual




Figure 4.2: Typical nonlinear activation functions. (a) Sigmoid, (b) Tanh, (c) ReLU, and (d) LeakyReLU. (Source: [9])

data under a variety of physical transformations [150]. They are conceptually similar to MLPs, but are translationally and rotationally invariant [151]. Additionally, CNNs introduce nonlinearity and pooling layers that adjust, limit, and summarize the generated outputs, allowing for robust and efficient pattern recognition, especially in image-based datasets [152]. Functionally, CNNs learn and dynamically adapt spatial hierarchies of features by passing an optimized feature extractor over smaller subsections of the data, summarizing the dominant features and associating them with ground truth labels that are then used to update weights and biases of the classification network. Figure 4.3 shows the general architecture of a CNN (Figure 2 from [10]). Here, input nodes are signals that are passed to the current layer, which may come from the output of hidden nodes in the previous layer or may be the input image itself. The filters perform a convolution with the signal from the input nodes. Pooling filters are low-pass convolutional filters that re-




Figure 4.3: CNN architecture example. An image undergoes two convolution and max pooling functions before the results are flattened, passed through a fully connected MLP, then classified into 2 categories at the output layer. (Source: [10])

duce dimensionality before the feature is extracted as output, enabling examination of the data at larger scales. Output nodes form a representative subset that is then flattened and passed through a feedforward MLP to categorize the image into one of two potential classes.

4.2 CNN for Detecting UXO in Radargrams

Machine learning techniques traditionally applied to GPR data include supervised learning frameworks such as CNNs and support vector machines (SVM) since UXO detection is a binary classification problem. Most studies rely on experimental data, where real or analog munitions casings are buried in a controlled environment and GPR surveys are conducted. The resultant data are then pre-processed using either open-source or proprietary software to produce radargrams, as described in Chapter 1. Reflection hyperbolas appear in radargrams at locations where the material has a different dielectric property than the surrounding soil matrix. Metal and plastic UXO have particularly large dielectric contrasts with common minefield soil matrices and generally produce strong reflection hyperbolas. Neural networks are then trained to detect and



Figure 4.4: Analog UXO Field Site. Pictures on the right show the analog UXO devices made of metal (top), plastic (middle), and ceramic (bottom).

classify these hyperbolas as either UXO or non-UXO. SVM algorithms rely on separating different classes with a hyperplane and predicting which side of the hyperplane new data points belong to. This technique has been moderately successful at identifying UXO in the few studies it has been applied to, but it constrains depth poorly [153]. CNN-based studies are showing excellent potential and can achieve high accuracy [154, 155]. Accuracy rates of 89-92% with false positive rates between 4-7% and false negative rates of 4% are achieved in cases where a sliding window technique is applied to an exhaustive dataset, such as in Nunez-Nieto's 2014 study [155]. We conducted a GPR study to explore the extent to which UXO detection is feasible using a CNN with data collected in suboptimal field conditions.

We constructed a 5 m x 5 m survey area in a densely vegetated region near College Park, MD, to match realistic minefield conditions in countries with high vegetation. Three UXO analogs (non-lethal objects that approximate UXO dimensions and material properties), made




Figure 4.5: Example pre-processed GPR radargrams from the UXO experiment. The horizontal axis of the radargram represents 5 m length, and the vertical axis represents 0-2 m depth. The upper panel contains UXO (location indicated with a red box and zoomed inset). The lower panel contains no UXO.

of (1) metal, (2) plastic, and (3) ceramic were buried at 15-20 cm depth along the diagonal of the field area. We chose a survey grid with 0.5 m spacing in North-South (N-S) and East-West (E-W) orientations. After measuring and marking the survey grid, we calibrated the GPR scroll wheel to ensure proper spatial measurements. N-S transects and E-W transects were collected using a lawnmower-style pattern, avoiding trees and fallen limbs as necessary. In total, 25 unique transects were collected throughout the experiment. The GPR data were then processed using standard practices outlined in [31] and radargrams were exported as JPEG images to be fed into the ML network. Figure 4.5a shows two examples of pre-processed data, one with UXO present (top, indicated with a red box and zoomed-in inset), and the other without (bottom).

To ensure a robust dataset for training the CNN, we performed data augmentation using Keras ImageDataGenerator [156], which takes an image and performs user-set alterations including horizontal shifts, vertical shifts, random zooms, flips, and shearing, among others. Horizontal shifts, horizontal flips, shearing, brightness variation, and random zoom alterations were applied

to each image until the total dataset reached 5000 images, 80% of which was assigned for training, and 20% for testing. The training dataset consisted of 2000 radargrams containing UXO and 2000 radargrams containing no UXO. The testing dataset consisted of 500 radargram images containing UXO and 500 radargram images containing no UXO. All images were resized from 1530 x 680 pixels to 90 x 40 pixels, maintaining the original aspect ratio.

We designed the CNN architecture in Anaconda Spyder [157], using Keras with a TensorFlow backend. The architecture consists of three convolutional layers with max-pooling layers after each convolution, followed by two fully connected dense layers, and a binary output layer that reports if the radargram image contains UXO or not. Figure 4.6 shows an illustration of the architecture. All max-pooling layers are 2 x 2 dimension. The first convolution layer contains 32 filters, with a kernel size of 5 x 5 and a max pooling layer. The second convolution layer contains 48 filters, with a kernel size of 4 x 4 and a max-pooling layer. The third convolution layer contains 48 filters, with a kernel size of 3 x 3 and a max-pooling layer. The output from convolution layer 3 is flattened, then fed into the first fully connected layer, which has 64 nodes. The second fully connected layer has 16 nodes, and is fed to the output layer, which has 1 node.

To ensure all testing and training data would be included when constructing the model, 4 epochs were assigned, with a batch size of 15. ReLu activation was used for all convolution layers and fully connected layers, while the output layer used sigmoid, as that is the more appropriate activation function for binary classification. Padding was applied to all convolution layers to keep the output size the same as the input size. Binary cross-entropy, which is commonly used for classification problems, was used as the model's loss function. The model optimizer selected was Nadam, which is essentially RMSprop with Nesterov momentum and was chosen to avoid both stalling and oscillating while training the model. The learning rate was set to 0.001, the




Figure 4.6: CNN Architecture for UXO Detection in Radargrams. Input images with 90 pixel width and 40 pixel height are fed through a series of convolution and pooling layers, followed by 2 fully connected layers before the output layer. The first convolution layer has 32 filters and a kernel size of 5 x 5. The second and third convolution layers have 48 filters, with 4 x 4 and 3 x 3 kernels, respectively. The first fully connected layer has 64 neurons, the second fully connected layer has 16 neurons, and the output layer has 1 neuron.

default rate for Nadam. To prevent overfitting, 0.1 dropout was applied to the fully connected layers only.

During initial training, RMSprop was used for optimization, and no regularization was applied to the hidden layers. RMSprop consistently stalled by epoch 2, so other optimizers were explored. Using Adam resulted in training accuracy oscillations, and the model ultimately underfitted. Once Nadam was applied, the training neither stalled nor oscillated, so it was chosen in the final design. Batch normalization was tested on all layers, which resulted in either underfitting or overfitting, no matter the combination of layers on which it was applied. Using batch normalization also decreased the average speed for each epoch, operating between 7-8 s/step. Dropout was tested in varying amounts on all layers, with the best performance at 0.1, and only when applied to the fully connected layers preceding the output layer. Dropout had a regularizing effect on training and allowed the model to generalize well. Adding dropout also increased the average speed for each epoch from 3s/step to 36 ms/step, allowing for quicker hyperparameter tuning and

evaluation. Batch sizes of 10, 15, 20, 30, and 50 were then explored. We found that a batch size of 10 resulted in an error message that the model ran out of data when testing, so it was discarded as an option. While batch sizes of 20, 30, and 50 produced similar results as a batch size of 15, they increased the average speed per epoch, and a batch size of 15 resulted in the lowest average speed and highest model performance. Increasing the number of epochs from 4 to 5, 10, 20, 25, and 30 overfit the data by epoch 5, so we chose 4 epochs for the final design. To evaluate if each layer was necessary for the model to learn (which has important computational efficiency implications), we tracked the weight histograms for each epoch, as seen in Figure 4.7. The horizontal axis represents the true weight applied to the model, and the vertical axis represents each epoch starting at 0. The shape and amplitude of the model weights change in every layer for every epoch, suggesting that each is making a significant contribution to the model's learning and should not be deleted from the CNN architecture. If the model weight curves appeared flat or unchanging between each epoch, it would suggest the layer may not be necessary for the model to learn effectively and should be discarded to maintain computational efficiency.

After fine-tuning the hyperparameters, we found that the CNN was able to achieve high training and testing accuracy, as well as low training and testing loss. Figure 4.8 shows examples of two input radargram images as they are transformed in each epoch. As the image is convolved, certain features are recognized by the CNN as important, such as localized dark spots, while others are recognized as unimportant and are eliminated by epoch 4.

Training accuracies between 89-99% and testing accuracies between 84-95% are seen in Figure 4.9, which displays the results of 10 runs of the CNN. While the exact rates vary with each run, the training and testing accuracy generally increases with each epoch. Training and testing loss decrease with each epoch, with exceptions in runs “grateful-haze-24” and “good-pond-14”,




Figure 4.7: Model weight histograms for each CNN layer. The distribution of true model weights is displayed at each epoch step, where epoch 0 is the first epoch, and epoch 3 is the fourth epoch. Changes in shape or amplitude between epochs suggest the model is using that layer to learn successfully. Changes in shape and amplitude are present for all layers, supporting the use of all in the final CNN design.




Figure 4.8: Examples of input radargram image changes after each epoch. (a) shows an example image containing UXO, while (b) does not contain UXO. Important features are kept by the CNN after each epoch, while unimportant features are discarded.

where testing loss somewhat increased between epochs 3 and 4. The consistent increase in testing accuracy and decrease in testing loss between each run suggests that the model is mostly stable, and the results are reproducible within an acceptable margin of error.

Figure 4.10 shows the confusion matrix and classification report for the CNN. The confusion matrix tracks how often the model classifies images correctly and incorrectly. A true positive is the case when the model correctly identifies that UXOs are present, and a true negative is the case when it correctly identifies that no UXOs are present. False positives represent times when the model flags the radargram for UXO, where none exists. The most dangerous case given the nature of the problem, however, is a false negative, where the model incorrectly classifies the radargram as “safe” when it *does* contain UXO. This CNN has a false negative rate of 4.90%, which outperforms similarly designed networks, though does not reach the goal of 2%, the standard set by the ordnance detection community. The F-score for a model is a weighted




Figure 4.9: Loss and accuracy for training and testing phases of the UXO CNN for 10 consecutive runs. Panels (a) and (b) show loss values for training and testing, respectively. While the rates vary, each run shows decreasing loss with increasing epoch. Panels (c) and (d) show accuracy values for training and testing, respectively. Both training and testing accuracies increase with increasing epoch, but the testing accuracy lags slightly behind training accuracy. This is likely a function of the model's weights and biases changing drastically in epochs 1 and 2. Maximum training accuracy values range between 86-99%, while testing accuracy values range between 84-95%.




Figure 4.10: UXO CNN Confusion Matrix and Classification Report. The upper panel shows the confusion matrix for the CNN, where 0 represents “No UXO Detected” and 1 represents “UXO Detected”. A false negative is the most dangerous case for a de-miner; the CNN was able to produce less than 5% false negatives for the 1000 radargrams used for validation. The F-1 score represents the weighted average of precision and accuracy for the model, where 1 is a perfect model.

representation of the precision and recall of a binary classification system. An F-score of 1 is considered to be perfect; this CNN has a score of 0.91, suggesting that it is strong, with some room for improvement.

The results from this work suggest that a CNN is an appropriate tool for identifying whether radargrams contain analog UXO or not, even when the data is collected in suboptimal field conditions. With only 4 training epochs and a small batch size of 15, the CNN was able to produce a maximum of 95% testing accuracy, which is higher than most of the accuracies achieved for similar experiments. Given the nature of the problem, however, 95% may not be sufficient for dealing with potentially explosive items. This CNN architecture and model are limited to classifying if UXO exists *somewhere* in the radargram image, without giving any information about the exact location or depth, which is a major limitation of the approach. The original radargram image set

was sparse and abundant data augmentation was used to generate a larger dataset; this should be taken into account when evaluating the performance of this CNN, as included radargrams may not be representative of real-world variability in hyperbola distribution and amplitude. The classification matrix confirms that the model is generalizing well, though it does not achieve a false negative rate of 2%, which would be ideal.

4.3 CNN for Magnetic Source Parameter Estimation

As we discussed in Chapter 3, detecting small signals associated with landmines, UXO, and ERW is a challenging task. Equally challenging is characterizing the source parameters associated with detected targets such as burial depth, magnetic moment, and orientation [158]. Modern approaches rely primarily on performing inversions of field data using a forward model that approximates the UXO as prolate spheroids, spheres, simple dipoles, or volumes of dipoles [159–163]. In this section, we examine the utility of multiple invariants derived from the FDMGT described in Chapter 3 and explore their merit as inputs to a CNN that estimates source parameters of small, detected ferrous targets.

While CNNs are used primarily for classification tasks, their architecture can be adapted to perform regression tasks. We estimate parameters such as magnetic moment and burial depth by adding a regression layer at the end of the CNN, using the weighted node values as inputs to the regression model. We designed the CNN using the same approach detailed in Section 4.3, tracking the model’s ability to learn with varying numbers of convolution layers, input nodes, and filter sizes, and then tuning hyperparameters to arrive at the most computationally efficient architecture. Figure 4.11 shows the CNN architecture of the network designed for magnetic source




Figure 4.11: CNN architecture for estimation source parameters using FDMGT invariant maps.

parameter estimation. It consists of an input layer, where 2-D FDMGT invariant maps are used as individual input channels, followed by four convolution layers with average pooling, batch normalization, and ReLU activation between convolution layers. The output from the fourth convolution layer is then flattened to a fully connected layer, followed by a regression layer that estimates the expected value of the parameter of interest.

One can use as many input channels as desired to explore the impact of different FDMGT invariants on the model's ability to estimate the source parameter of interest, but it should be noted that more inputs do not necessarily guarantee better results given the computational cost. For our purposes, we restricted the number of input layers to a maximum of five to better understand the information gained from different invariant combinations.

We generated 12,314 instances of synthetic data for a buried dipole calculated using TetraMag's sensor geometry, varying the orientation of TetraMag, magnetic inclination and declination values, magnetic moment values, burial depths, and location of the target in a 5 m x 5 m domain. We then calculated the FDMGT and corresponding invariants using the methods described in Chapter 3 and organized invariant maps of interest as input layer channels for the CNN. Data

were split into 70% training, 20% validation, and 10% testing sets. This ensured that the testing data were not "seen" by the CNN while the model was being constructed and trained.

To estimate the magnetic moment of the dipole, we systematically tested both traditionally measured magnetic data (B_t , B_x , B_y , B_z) as well as the magnetic gradient tensor-derived invariants as input channels to the CNN. Figure 4.13 shows a summary of the CNN performance for each method in terms of accuracy within four error tolerances (within 0.06, 0.05, 0.04, and 0.03 Am²). Instances when the CNN was trained with a single channel are labelled with only the data type used in that channel. Instances where multiple input channels were used list the included data types, separated by commas. When only B_t is used as an input (which any total field magnetometer would provide), magnetic moment is estimated to within 0.03 Am² with 25% accuracy. When vector component information is available (which a single vector magnetometer would provide) and B_x , B_y , and B_z are included as three input channels, the magnetic moment within 0.03 Am² increases to 65% accuracy. The performance of single-channel magnetic tensor invariants such as eigenvalues λ_1 , λ_2 , and λ_3 within 0.03 Am² is typically < 50%, but when invariants are combined, the model performs dramatically better. When using $|G_S|$, $\|G_S\|_F$, λ_1 , λ_2 , and λ_3 invariants as input layers, for example, the magnetic moment is estimated within 0.03 Am² at > 95% accuracy. Figure 4.12 shows the true magnetic moment vs. the magnetic moment predicted by the trained CNN in this instance. The mean orientation angle error for this model is 1.5°. Accurately estimating the magnetic moment within these error bounds amounts to a full order-of-magnitude improvement on methods that are currently being used such as those in [81, 160, 164]. This highlights that the full magnetic gradient tensor contains crucial information about small-scale targets that are not sufficiently captured by measuring B_t alone, or even B_x , B_y , and B_z .




Figure 4.12: Magnetic moment estimation from source parameter CNN. The red line is a 1:1 reference line.

To estimate the burial depth of the dipole, we followed a similar approach as described above to evaluate the CNN performance of traditionally measured magnetic data and invariants of the gradient tensor. Figure 4.15 shows a summary of the CNN performance for each method in terms of accuracy within four error tolerances (within 4, 3, 2, and 1 cm). The total field, B_t , performs the worst, estimating depth within 3 cm at $< 30\%$ accuracy. Similar to the CNN designed to estimate magnetic moment, including the three vector components B_x , B_y , and B_z increases the depth prediction accuracy to 55% within 3 cm, but pales in comparison to the performance of magnetic gradient tensor invariants, especially when used in combination. Figure 4.14 shows the resulting true burial depth vs the burial depth predicted by the trained CNN when the ratio $\|G_A\|_F/\|G_S\|_F$, and eigenvalues λ_1 , λ_2 , and λ_3 are used as input channels. The burial depth of the target can be estimated to be within 4 cm tolerance at 99.43% accuracy, and within 3




Figure 4.13: Magnetic data of interest as included CNN layers vs. magnetic moment estimation accuracy.




Figure 4.14: Burial depth estimation from CNN. The red line is a 1:1 reference line.

cm at 94.4% accuracy. If one wishes to estimate to within 1 cm, the accuracy falls to 78.43% but outperforms all other invariant combinations that were tested. This confirms the depth resolving power of both the eigenvalues of FDMGT, and the $\|G_A\|_F / \|G_S\|_F$ ratio described in Chapter 3, especially when compared to traditionally measured magnetic data. Being able to estimate the burial depth at cm scale is sufficient given current demining practices, and amounts to a modest improvement on current depth estimation techniques such as [165], but has the benefit of avoiding the upward continuation calculations required for the depth from extreme points (DEXP) or multi-ridge methods.




Figure 4.15: Magnetic data of interest as included CNN layers vs. burial depth estimation accuracy.

4.4 Conclusions

This study demonstrates how machine learning techniques such as MLPs and CNNs are useful for both detecting small-scale, small-amplitude targets and estimating their source parameters. When GPR data is available, the CNN from Section 4.2 may be used to highlight radargrams of interest that likely contain UXO, with 95% accuracy. The addition of a sliding window to automatically identify the location of the detected UXO hyperbola would make this method even more useful for rapidly identifying regions of interest. When a regression layer is added to the end of the CNN architecture, source parameters such as magnetic moment and burial depth may be estimated from FDMGT invariant maps with high accuracy and acceptable error margins. The CNN’s ability to estimate both magnetic moment and burial depth highlights that important information is contained within the FDMGT that is not necessarily captured by measuring just B_t , or the vector components B_x , B_y , and B_z , especially at these small scales. The inclusion of ML as a data processing tool enables the fusion of multiple geophysical modalities when detecting and characterizing targets in the shallow critical zone, be they anthropogenic or geological in origin.

Chapter 5: Method of Self-Noise Reduction for Drone-Based Magnetometry

In this chapter, we present a method we have developed for modeling and reducing unwanted magnetic signals generated by vehicle self-noise when conducting drone-based surveys. We begin by motivating the need to model vehicle self-noise when using rigidly mounted magnetometers and discuss their sources. We then describe our workflow and the construction of our model. Finally, we validate our model using field data and discuss the strengths and limitations of our technique.

5.1 Introduction

Detecting small-scale, small-amplitude anomalies is particularly challenging, whether using vector or scalar magnetometers. This is partly due to cultural and space weather noise sources that can obfuscate anomalies when used in isolation [159]. Magnetic gradiometry largely solves the issue of cultural and space weather noise [166, 167], but does not address the impact of the drone platform, which will produce a non-trivial, time-varying magnetic signature as part of the total measured signal. Magnetometer systems are often slung-loaded to avoid these unwanted signals. However, swinging motions during flight restrict users to total field magnetometers due to the pointing sensitivity of vector magnetometers [168]. Additionally, slung-loaded systems are inappropriate in cases requiring very low-altitude flying. Total field magnetometers provide

less information than triaxial vector magnetometers, limiting the ability to detect and constrain dimensions of small-scale anomalies. To mitigate space weather and cultural noise sources and maximize the amount of information collected in the field, we utilize magnetic gradiometry with an array of rigidly-mounted triaxial fluxgate sensors, called TetraMag, aboard a Freefly Alta X drone, capable of carrying 15 kg payloads. This necessitates characterizing the combined magnetic signature of the drone to remove from magnetic data.

5.1.1 Sources of Magnetic Self-Noise

Magnetic self-noise sources of drone platforms include current draw from DC motors, motions of permanent magnets inside the motors, vibration, vehicle attitude, and any payload instrumentation being flown. Figure 5.1 illustrates these noise sources on the Freefly Alta X platform. Owing to Ampere’s Law, DC brushless motors (BLDC) aboard the Alta X contribute magnetic signature as current is drawn from the DC power bus to each motor (labeled *a* on Fig. 5.1) [169–171]. The motion of permanent magnets inside BLDC motors produces time-varying magnetic fields, and any conductive drone components moving through Earth’s field will also produce magnetic fields, resulting in additional self-noise [172], with label *d*. Additionally, small-scale motions from the drone platform vibrating can be faster than the ability for onboard sensors to measure and compensate for, creating noise in collected data (label *c*) [37]. Payload instrumentation, labeled *e* on Fig. 5.1, may also contribute magnetic noise depending on their specifications. Triaxial fluxgate magnetometers measure along the axis of each sensor core, so Earth’s ambient field is sampled as the projection of that field onto the axis of each component (label *b*). For a rigidly mounted system, microradian uncertainty in sensor orientation and posi-




Figure 5.1: Sources of magnetic noise on a drone. (a): current draw from DC power bus to motors. (b): Earth’s ambient magnetic field. (c): vibration of the platform. (d) motion of permanent magnets inside BLDC motors. (e): payload instrumentation.

tioning with respect to the attitude of the drone (yaw, pitch, and roll) will result in the creation of spurious magnetic signals or the suppression of true magnetic signals, contributing to the total vehicle self-noise.

This motivates the need to model the combined magnetic signature of the drone to cancel out their contribution to measured magnetic data for the rigidly-mounted triaxial fluxgate magnetic gradiometer.

5.2 Magnetic Noise Reduction Model

Figure 5.2 shows the workflow for the magnetic noise reduction model and the creation of anomaly detection maps. TetraMag data are collected using a Raspberry Pi 4 equipped with a real-time clock, while the Alta X logs its telemetry data using its embedded Black Cube, with timestamps from real-time kinematic (RTK) GNSS data. We begin with data preparation (described in Section 5.3.2), including calibration of TetraMag data, synchronization with the




Figure 5.2: Magnetic noise reduction workflow.

telemetry data, and downsampling to 4 Hz. We leverage telemetry data the Alta X measures to construct the TetraMag Predictive Model. Since we are only interested in the total contribution from the combined magnetic noise sources, multiple potentially useful parameters could help capture information about the combined signal, including:

- GPS Position: latitude (λ), longitude (τ)
- Altimeter Elevation (h)

- Vehicle Attitude: yaw (ψ), pitch (θ), roll (ϕ)
- Motor Current Draw ($\vec{I}$)
- Motor Speed: RPM ($\vec{\Omega}$)
- Etc...

We train candidate linear regression models using parameters from the telemetry dataset described below.

5.2.1 Expected $\vec{B}$ Value

We combine the latitude, longitude, elevation, and attitude information with the known TetraMag sensor locations to calculate expected B_x , B_y , and B_z values for each sensor. We begin by querying the world magnetic model (WMM) for the expected Earth B_x , B_y , and B_z values for the drone’s measured latitude, longitude, and elevation from an externally integrated laser rangefinder. Each sensor’s position is corrected for the known separation between the drone’s attitude sensor before querying the WMM for that sensor. Due to ground effects impacting barometer stability, the laser rangefinder provides more reliable altitude measurements for our very low-altitude (< 1 m) flights than the Alta X’s combination IST8310 magnetometer and BMP388 barometer. We then use the telemetry-logged yaw, pitch, and roll to construct the cosine rotation matrix and project the ambient expected field onto the plane of the drone. We abbreviate these estimations as "Exp" in the model construction.

5.2.2 Motor Current Draw, Speed And Their Time Derivatives

DC motors convert current and voltage into mechanical motion (+ heat). They are generally governed by two equations: the time rate of change of the conservation of energy and the motor's voltage balance equation. The conservation of energy equation can be written as:

$$IV = I^2R + \tau\Omega \quad (5.1)$$

where I is current, V is the voltage applied to the motor, R is the resistance that comes from the motor's armature, τ is torque, and Ω is the angular speed [173]. The motor voltage balance equation is given as follows:

$$V = IR + L\dot{I} + V_{EMF} \quad (5.2)$$

where L is motor inductance, $\dot{I}$ is the time derivative of current, and V_{EMF} is the back electromotive force (EMF), which is associated with the work the DC motor does to convert either electrical to mechanical energy, or mechanical to electrical energy.

Motor current draw, $\vec{I} = (I_1, I_2, I_3, I_4)$, is measured for the four motors aboard the Alta X via its built-in telemetry system. Each motor is wired to a central power distribution unit, so we can expect a magnetic field around each of these wires as the motor draws current from the central batteries to approximately follow the Biot-Savart Law [174]:

$$\vec{B}(\vec{r}) = \frac{\mu_0}{4\pi} \int_C \frac{Id\vec{l} \times \vec{r}'}{|\vec{r}'|^3} \quad (5.3)$$

where C is the path, $d\vec{l}$ is the differential length along the path in the direction of the current flow, $\vec{r}' = \vec{r} - \vec{l}$ is the full displacement vector from the measurement point $\vec{r}$ to the point along the wire, and μ_0 is the permeability of free space. By including $\vec{I}$, we consider the conservation of energy of the motors and torque since the current is directly proportional to torque [175]. We include the time derivative of current, $\dot{\vec{I}}$, to account for the motor's voltage balance. We also include the Hilbert Transform of each motor's current measurement, $h\vec{I}$, the imaginary part of the analytic signal of $\vec{I}$, to capture information about any potential phase delays.

Similar to the current draw, the motor speed, $\vec{\Omega} = (\Omega_1, \Omega_2, \Omega_3, \Omega_4)$ is measured for the four motors by the Alta X telemetry system. A motor may be thought of as being in series with a resistor [176]. As the induced current increases, V_{EMF} will also increase, proportional to Ω (see Eqn 5.1 and 5.2). To account for both the back EMF and the conservation of energy, we include $\vec{\Omega}$ in our model. Since the drone will not necessarily be moving at a constant velocity, we also include the time derivative of motor speed, $\dot{\vec{\Omega}}$, which may account for eddy currents in the motors that arise by Faraday induction [177]. Additionally, we include the Hilbert Transform of each motor's speed, $h\vec{\Omega}$, to ensure potential phase delay information is considered.

5.2.2.1 Candidate Models

We construct three separate models to explore the extent to which these telemetry and telemetry-derived parameters are useful in capturing magnetic noise from the drone:

- Model A: Expected $\vec{B}, \vec{I}, \vec{\Omega}, h\vec{I}, h\vec{\Omega}$
- Model B: Expected $\vec{B}, \dot{\vec{I}}, \dot{\vec{\Omega}}, h\vec{I}, h\vec{\Omega}$
- Model C: Expected $\vec{B}, \vec{I}, \dot{\vec{I}}, \vec{\Omega}, \dot{\vec{\Omega}}, h\vec{I}, h\vec{\Omega}$

5.3 Field Data

This section describes our field area and the data processing and preparation steps required to train the magnetic self-noise model.

5.3.1 Inert Minefield

Our site, located near Glenn Dale, MD, lies between three fields primarily used for agromomic crop research and is in a soil map named for the Downer and Hammonton soil series (Coarse-loamy, siliceous, semiactive, mesic Typic Hapludults, and Coarse-loamy, siliceous, semiactive, mesic Aquic Hapludults). We assembled a 17.5 m x 7.5 m survey area using a collection of inert landmines and cluster munitions at realistic depths for their typical deployment. Figure 5.3 shows the location of the field site (top left), a picture of the Freefly Alta X with our multi-sensor payload (top right), as well as an illustration of the survey layout, with row labels A, B, and C, and column labels 1-7. Targets known as ferrous are in red, possibly ferrous are in yellow, non-ferrous are in white, and blank control holes are in black. Table 5.1 contains a complete description of each target, with dimensions and burial depths.

5.3.2 Data Preparation

We flew a clover-shaped mission 8 m above the survey area to calibrate our magnetic data and applied the least-squares ellipsoid fitting calibration process described in [43, 126, 129]. We then flew a grid pattern survey (N-S and E-W orientations) at 0.5 m altitude with 30 cm spacing between transects across the field area at 0.5 m/s. We synchronized TetraMag measurements with telemetry data by cross-correlating Alta X’s onboard magnetometer with the closest mag-

Target Label	Type	Dimensions	Burial Depth
A1	Control hole	N/A	N/A
A2	AT Mine	24 cm D, 12 cm H	7 cm
A3	Dummy Bomb	33 cm L	10 cm
A4	IED Analog	4 cm D, 10 cm L	13 cm
A5	EM Mine	30 cm D, 10 cm H	5 cm
A6	Dummy Grenade	5 cm D, 8 cm H	8 cm
A7	Simulant AT mine	30 cm D, 10 cm H	10 cm
B1	M42 Submunition	3 cm D, 5 cm H	10 cm
B2	AT mine	15 cm D, 5 cm H	10 cm
B3	PDM AP mine	6 cm L, 10 cm W	10 cm
B4	Dummy Bomb	33 cm L	10 cm
B5	Claymore AP	22 cm L, 4 cm W, 12 cm H	5 cm
B6	Simulant AT mine	12 cm D, 4 cm H	7 cm
B7	Bouncing Betty AP mine	10 cm D, 15 cm H	12 cm - 0 cm
C1	AT mine	24 cm D, 8 cm H	9 cm
C2	Control hole	N/A	N/A
C3	Metal AT mine	33 cm D, 13 cm H	10 cm
C4	BLU-26 Bomb	6 cm D (spheroid)	10 cm
C5	Artillery Shell Analog	33 cm L, 7 cm H	5 cm
C6	Plastic AP mine	10 cm D, 5 cm H	10 cm
C7	Simulant AT mine	30 cm D, 8 cm H	8 cm

Table 5.1: Target specifications for Glenn Dale, MD, field site. Abbreviations: D (diameter), H (height), L (length), W (width), AT (anti-tank), AP (anti-personnel), IED (improvised explosive device), EM (electromagnetic), PDM (pursuit deterrent munition).




Figure 5.3: Inert Minefield and target locations near Glenn Dale, MD. The photo inset shows the Alta X drone with our geophysical sensor payload.

netometer to the platform (magnetometer 4). We then downsampled all telemetry data to match TetraMag’s 4 Hz sampling frequency. To evaluate the extent to which filtering is required to suppress vibration noise from the drone, we compared the unfiltered data to a 0.25 Hz low-pass filtered version. Figure 5.4 shows anomalies, defined as the measured $\vec{B}$ value subtracted from the Expected $\vec{B}$ value, for magnetometer 1, z-component (B_{1z}) for the unfiltered case (top) and filtered case (bottom). The spatial coherence increases dramatically once the 0.25 Hz low-pass filter is applied, suggesting that vibration noise contributes significantly to the measured data and should be filtered out, especially when small-scale, small-amplitude anomalies are expected. As a result, we applied the 0.25 Hz low-pass filter to both TetraMag and telemetry data before constructing candidate models.




Figure 5.4: Comparison between no filter and 0.25 Hz low-pass filtered data. The filtered data has better spatial coherence as it suppresses vibration from the drone flying.

5.4 Results

Here, we present the results for candidate models A, B, and C using data collected at the inert minefield. We begin by discussing the relative importance of each model's included parameters by normalizing the input variables and data to a standard deviation of 1 and comparing the resultant coefficients. We then show each model's performance with respect to root mean squared error (RMSE) and spatial coherence.

5.4.1 Candidate Model A Coefficients

Figure 5.5 shows how Model A's coefficient amplitudes of included parameters compare, component-by-component, for each magnetometer. Expected $\vec{B}$ values are omitted since they are included in all models. Magnetometers 1 – 4 ($B_1 - B_4$) are displayed as "Mag1", "Mag2",

”Mag3”, and ”Mag4”, respectively. Model A may be written as:

$$B_{Exp}^{\vec{}} + \begin{bmatrix} \vec{a}_1 & \vec{b}_1 & \vec{c}_1 & \vec{d}_1 \\ \vec{a}_2 & \vec{b}_2 & \vec{c}_2 & \vec{d}_2 \\ \vec{a}_3 & \vec{b}_3 & \vec{c}_3 & \vec{d}_3 \\ \vec{a}_4 & \vec{b}_4 & \vec{c}_4 & \vec{d}_4 \end{bmatrix} \begin{bmatrix} \vec{I} \\ \vec{\Omega} \\ h\vec{I} \\ h\vec{\Omega} \end{bmatrix} = \begin{bmatrix} {}^A\vec{B}_1 \\ {}^A\vec{B}_2 \\ {}^A\vec{B}_3 \\ {}^A\vec{B}_4 \end{bmatrix} \quad (5.4)$$

where ${}^AB_{1-4}$ is Model A’s predicted magnetic field for each magnetometer, $B_{Exp}^{\vec{}}$ is the set of expected $\vec{B}$ values for each magnetometer, $\vec{I}$, $\vec{\Omega}$, $h\vec{I}$, and $h\vec{\Omega}$ are the input variables, and $a_{(1-4)}$ - $d_{(1-4)}$ represent the model coefficients that scale each variable. When modeling $B_{(1-4)x}$, the current draw from motor 2 (I_2) dominates the model coefficients for all magnetometers. Ω_1 contributes the least for B_{1x} , but is the second most dominant for B_{1y} . Hilbert transforms of current draw and RPM make significant contributions across all components, suggesting the existence of nonlinearity that can be approximated as a phase shift. The relative importance of each parameter is consistent across all magnetometers in $B_{(1-4)x}$ and $B_{(1-4)y}$. Still, parameters for B_{4z} (located closest to the drone body) are different, with Hilbert transforms of both Ω ($h\Omega$) and I (hI) being the most prominent. I_3 does not contribute much to $B_{(2-4)y}$ for magnetometers, but does for magnetometer 1.

5.4.2 Candidate Model B Coefficients

Figure 5.6 shows the coefficients for variables in Model B. Model B may be written as:




Figure 5.5: Model A coefficients. Input data and model variables have been normalized to $\text{std} = 1$. The four magnetometers of TetraMag are labeled "Mag1", "Mag2", "Mag3", "Mag4".

$$B_{Exp}^{\vec{}} + \begin{bmatrix} \vec{a}_1 & \vec{b}_1 & \vec{c}_1 & \vec{d}_1 \\ \vec{a}_2 & \vec{b}_2 & \vec{c}_2 & \vec{d}_2 \\ \vec{a}_3 & \vec{b}_3 & \vec{c}_3 & \vec{d}_3 \\ \vec{a}_4 & \vec{b}_4 & \vec{c}_4 & \vec{d}_4 \end{bmatrix} \begin{bmatrix} \dot{\vec{I}} \\ \dot{\vec{\Omega}} \\ h\vec{I} \\ h\vec{\Omega} \end{bmatrix} = \begin{bmatrix} {}^B\vec{B}_1 \\ {}^B\vec{B}_2 \\ {}^B\vec{B}_3 \\ {}^B\vec{B}_4 \end{bmatrix} \quad (5.5)$$

where ${}^B B_{1-4}$ is Model B's predicted magnetic field for each magnetometer, $B_{Exp}^{\vec{}}$ is the set of expected $\vec{B}$ values for each magnetometer, $\dot{\vec{I}}$, $\dot{\vec{\Omega}}$, $h\vec{I}$, and $h\vec{\Omega}$ are the input variables, and $a_{(1-4)}$ - $d_{(1-4)}$ represent the model coefficients that scale each variable. Hilbert transforms of I and Ω dominate across all components of all magnetometers, suggesting increased nonlinearity when only the gradients of I and Ω are included in the model. The amplitudes of all variable coefficients are larger than those from Model A. The relative importance of each variable is similar for $B_{1(x,y,z)}$, $B_{2(x,y,z)}$, and $B_{3(x,y,z)}$, but $\dot{I}_4$ and $\dot{\Omega}_4$ are more prominent in B_{4z} , and $h\Omega_4$ and $h\Omega_1$ are more prominent when predicting B_{4y} . $\dot{\Omega}_4$ contributes very little for $B_{1(x,z)}$, $B_{2(x,z)}$ and $B_{3(x,z)}$, but contributes significantly for the $B_{(1-4)y}$ components. $\dot{I}_1$ is consistently larger than $\dot{I}_2$ for all components and all magnetometers.

5.4.3 Candidate Model C Coefficients

Figure 5.7 shows the coefficients for variables in Model C. Model C may be written as:




Figure 5.6: Model B coefficients. Input data and model variables have been normalized to $\text{std} = 1$. The four magnetometers of TetraMag are labeled "Mag1", "Mag2", "Mag3", "Mag4".

$$B_{Exp}^{\vec{}} + \begin{bmatrix} \vec{a}_1 & \vec{b}_1 & \vec{c}_1 & \vec{d}_1 & \vec{e}_1 & \vec{f}_1 \\ \vec{a}_2 & \vec{b}_2 & \vec{c}_2 & \vec{d}_2 & \vec{e}_2 & \vec{f}_2 \\ \vec{a}_3 & \vec{b}_3 & \vec{c}_3 & \vec{d}_3 & \vec{e}_3 & \vec{f}_3 \\ \vec{a}_4 & \vec{b}_4 & \vec{c}_4 & \vec{d}_4 & \vec{e}_4 & \vec{f}_4 \end{bmatrix} \begin{bmatrix} \vec{I} \\ \vec{\Omega} \\ \dot{\vec{I}} \\ \dot{\vec{\Omega}} \\ h\vec{I} \\ h\vec{\Omega} \end{bmatrix} = \begin{bmatrix} {}^C\vec{B}_1 \\ {}^C\vec{B}_2 \\ {}^C\vec{B}_3 \\ {}^C\vec{B}_4 \end{bmatrix} \quad (5.6)$$

where ${}^C B_{(1-4)}$ is Model C's predicted magnetic field for each magnetometer, $B_{Exp}^{\vec{}}$ is the set of expected $\vec{B}$ values for each magnetometer, $\vec{I}$, $\vec{\Omega}$, $\dot{\vec{I}}$, $\dot{\vec{\Omega}}$, $h\vec{I}$, and $h\vec{\Omega}$ are the input variables, and $a_{(1-4)}$ - $f_{(1-4)}$ represent the model coefficients that scale each variable. For $B_{(1-4)x}$, $h\Omega_1$ and Ω_2 dominate, while hI_4 and I_3 dominate for $B_{(1-4)z}$. $\dot{I}_3$ does not contribute significantly to $B_{(2-4)x}$, but does across all magnetometers for y- and z-components. Like Model A and Model B, magnetometer 4 has different dominant variables in the y- and z- components (Ω_2 , I_4 , I_3 , and I_1 for B_{4y} and $h\Omega_2$, I_1 , Ω_3 , and $h\Omega_3$ for B_{4z}).

5.4.4 Model Performance Comparison

Figure 5.8 summarizes each model's performance by the error between the model's predicted value of B and the TetraMag measured values (RMSE). Figure 5.8(i) shows RMSE by component for the models A, B, and C, and when only B_{Exp} is included. The largest errors occur for the B_{Exp} case, which only accounts for the drone's orientation with respect to the world magnetic model. Model B has consistently higher errors than models A and C, and model C has the lowest errors across all components. Z-component predictions produce the lowest RMSE values for all models, with ${}^C B_{4z}$ producing average < 60 nT errors. Figure 5.8(ii) shows measured B_{4z}




Figure 5.7: Model C coefficients. Input data and model variables have been normalized to $\text{std} = 1$. The four magnetometers of TetraMag are labeled "Mag1", "Mag2", "Mag3", "Mag4".

Component	B_{Exp} Only	Model A	Model B	Model C
B_x	1	2	2	2
B_y	< 1	1	1	1
B_z	2	3	3	3

Table 5.2: Average order of magnitude variance reduction by component for each candidate model.

(blue) and predicted ${}^C B_{4z}$ (dashed black) for one of the 0.5 m altitude survey flights, with the resulting errors (anomalies) in the bottom panel (iii). Model C recovers the longer periods and the majority of shorter periods from the measured signal. The remaining errors (anomalies) match typical landmine, UXO, and ERW anomaly ranges. Figure 5.8(iv) is a scatterplot of the predicted ${}^C B_{4z}$ vs measured B_{4z} values seen in panel (ii), and the red line is a 1:1 ($y=x$) reference. Model C predicted values closely follow the 1:1 line, amounting to three orders of magnitude variance reduction. Table 5.2 summarizes the remaining variance reduction estimates. While models A, B, and C all produce the same order of magnitude variance reduction, the smallest values for individual magnetometer components are in model C.

Table 5.3 provides a summary of the root mean squared error between different components of the finite-difference magnetic gradient tensor, $\mathbf{G}$, using both uncorrected and corrected TetraMag data for each model. Model C has the smallest RMSE for all components of $\mathbf{G}$, and models A, B, and C all perform better than the B_{Exp} correction.

Figure 5.9 shows the average B_x , B_y , and B_z anomaly maps for the inert minefield with the known target locations overlaid. We use basis spline interpolation to visualize the data, constructing a piecewise polynomial function that passes through each data point while minimizing error or curvature. We used 0.25 m knot spacing, L1 misfit to quantify the error, and a damping factor of 100 for all maps. The top row shows the average component maps for orientationally




Figure 5.8: Model performance summary. (i) shows the root mean squared error for each component by model. "Exp B" uses only the geometrically-compensated values. (ii) shows an example signal from Magnetometer 4, z-component ($^C B_{4z}$), with the values predicted by Model C. (iii) shows the anomalies from Model C. These values match the expected anomaly ranges for landmines, UXO, and ERW. (iv) shows predicted B_{4z} vs measured B_{4z} . The red line is a 1:1 reference.

G Component	B_{Exp} Only	Model A	Model B	Model C
B_{xx}	1880.37	1871.77	1871.87	1871.45
B_{yx}	209.93	161.84	168.53	159.99
B_{zx}	1374.23	1364.86	1364.99	1363.82
B_{xy}	726.79	712.75	712.69	711.40
B_{yy}	1398.341	1389.57	1389.22	1389.14
B_{zy}	219.97	106.64	109.58	106.40
B_{xz}	1611.19	1607.59	1607.45	1607.38
B_{yz}	1961.91	1960.60	1960.61	1960.58
B_{zz}	3259.63	3249.01	3249.14	3248.96

Table 5.3: Root mean squared error (RMSE) between different components of G using both uncorrected and corrected TetraMag data for each model.

compensated B values (B_{Exp}), and rows 2, 3, and 4 show maps for models A, B, and C, respectively. B_x , B_y , and B_z maps show very little spatial coherence for the B_{Exp} case but increase dramatically with models A, B, and C. Model C shows the strongest spatial coherence across all components. However, model A highlights many of the same features as model C, especially in the x-component. B_z anomalies are smaller than those for B_x and B_y , but the increase in spatial coherence is still dramatic for models A and C, highlighting areas with small-scale, small-amplitude signals. Locations of targets that are ferrous or suspected to be ferrous have black crosses, and non-ferrous or control hole locations have black circles. The increase in spatial coherence for each average component approximately localizes many of the targets with a spatial offset. Still, many of the signals blend together, making it difficult to differentiate them.




Figure 5.9: Average B_x , B_y , and B_z anomalies by magnetic noise model, with target locations indicated. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.

5.5 Discussion

Here, we discuss the results of our magnetic self-noise models and their implications. We begin by describing the impact of drone magnetic noise on data fidelity and then discuss the limitations resulting from TetraMag’s low sampling frequency.

5.5.1 The Impact of Drone Magnetic Noise on Data Fidelity

As one would expect, the largest contributing signal measured by magnetometers comes from Earth’s ambient field. For total field magnetometers, this may be partially compensated by using a base station magnetometer. Still, triaxial magnetometers are sensitive to orientation, so a base station alone is not sufficient for correcting triaxial data, especially on drones. As described in Section 5.2.1, we can calculate more accurate $\vec{B}$ correction values ($B_{Exp}^{\vec{}}$) by projecting the WMM ambient field vectors onto the plane of the drone using a direction cosine matrix derived from the telemetry’s logged yaw, pitch, and roll data. Figure 5.8(i) and Table 5.2 show that this results in 1, < 1 , and 2 orders of magnitude variance reduction for B_x , B_y , and B_z , respectively. While this may be sufficient for detecting large-scale anomalies, it is far larger than the expected signal amplitudes from small-scale anomalies such as landmines, UXO, and ERW. Including the logged current draw, motor speed, as well as their gradients and Hilbert transforms, in a linear model allows us to capture information about the drone’s generated magnetic self-noise, while a low-pass filter helps suppress some of the signals associated with vibration during flight. Models A, B, and C all produce the same order of magnitude variance reduction (see Table 5.2), but Model C produces the lowest RMSE overall. This suggests that the magnetic signals generated by the drone have both linear and nonlinear factors but that nonlinear factors are well-approximated

as a phase shift from the Hilbert transforms of $\vec{I}$, $\vec{\Omega}$, and their gradients.

The average anomaly maps in Figure 5.9 are similar to what one might expect from using a single triaxial fluxgate magnetometer. While the spatial coherence increases and noise is reduced using these models, the locations of ferrous targets remain uncertain, reinforcing the importance of leveraging invariants derived from the magnetic gradient tensor. Figure 5.10 shows one of the invariants of interest, the determinant of the symmetric part of the magnetic gradient tensor ($|G_S|$) for all models, with the known target locations overlaid. Many spurious signals are present when only B_{Exp} is included, which obfuscates target locations. Models A and B dramatically improve the resolution for different sections of the survey area. Model C, which incorporates elements from A and B, produces the highest-resolution map, localizing most targets. Chapter 6 introduces additional invariant maps based on the gradient tensor results computed from model C.

5.5.2 Sampling Frequency Limitations

The drone telemetry sensors have a sampling frequency of 25 Hz, while the MEDA FVM-400 sensors used by TetraMag in continuous data streaming mode are 4 Hz. While downsampling the telemetry data to match TetraMag prioritizes the retention of all magnetic data without the risk of introducing spurious signals, the loss of information from telemetry impacts both the maximum surveying velocity and the accuracy of yaw, pitch, and roll measurements included in the model. High data density is important for all magnetic surveys, but especially so when one expects the presence of small-scale, small-amplitude anomalies. To ensure proper data coverage for the scales of our known buried targets, we were limited to a flying velocity of 0.5 m/s. Flying slowly has the benefit of increasing the drone's stability at the limiting cost of battery life. Our




Figure 5.10: The determinant of the symmetric part of the full magnetic gradient tensor ($|G_S|$) for all models. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.

payload weight amounts to roughly 12 – 15 minutes of flying time for each set of batteries, which take upwards of 40 minutes to charge once used. Additionally, since the model relies on yaw, pitch, and roll measurements to correct the orientation of the drone, this information loss may limit our ability to properly describe the drone’s position at every timestep. This directly impacts the efficacy of the B_{Exp} calculation.

Our magnetic self-noise model for the drone outperforms current compensation techniques and enables the rigid mounting of sensitive triaxial magnetometers to the drone body. Still, it would benefit from a higher sampling magnetometer array. If TetraMag could sample at the same 25 Hz as the telemetry, we could fly more than six times faster for the same survey data density. This would obviate the need for downsampling, allowing for all yaw, pitch, and roll measurements to be included, which may improve the self-noise model even more. This, in turn, would increase overall data fidelity.

5.6 Conclusions

This study demonstrates how data logged by UAV telemetry may be leveraged to model and reduce vehicle magnetic self-noise. The method is useful for any drone-mounted magnetic system with an accessible telemetry log. It is standard for commercial drones to log yaw, pitch, roll, Ω , and battery status ($\vec{I}$), so no additional sensors are required to utilize the developed technique. However, if desired, this does not preclude using extra, higher-resolution accelerometers and/or current sensors to measure these parameters. To measure $\vec{I}$ in a non-invasive way, for example, one could attach split-core current sensors at each motor wire connection and log the data using either the flight controller or an external microcontroller fitted with a real-time clock.

Data fidelity and spatial coherence improve dramatically for models A, B, and C compared to results using $B_{Exp}^{\rightarrow}$ only. Model C, which includes currents, motor speed, and their derivatives, ultimately outperforms the others. It is consistent with the conservation of energy equation (5.1) and the motor voltage balance equation (5.2), which govern DC motors. Measurable magnetic signals from the vibration of the drone may be suppressed effectively with a low-pass filter, further increasing data fidelity. This method may be improved by using higher sampling frequency magnetometers, enabling faster surveying velocity and wider coverage per set of drone batteries. For the Freefly Alta X, a higher sampling frequency for TetraMag would ensure no telemetry data are lost while downsampling. Since the Alta X has a vibration-damping ring between its body and our custom landing gear, it is worth investigating the extent to which yaw, pitch, and roll measured by the drone matches the movement of TetraMag. Integrating an additional accelerometer on the landing gear would verify the accuracy of the telemetry measurements and inform future versions of the drone package. This method is not restricted to UAV-based magnetometry and may be applied to any vehicle that utilizes DC motors. Ground robotics systems such as the AeroVironment tEODor EVO [178], for example, which disarms and disposes of explosive ordnance, could be fitted with TetraMag and use these modeling methods to model and subtract out the vehicle's self-noise for UXO detection prior to removal.

Chapter 6: Landmine and ERW Detection using UAV-Based Multi-Modal Geophysics

In this chapter, we present the individual sensor and multi-modal results of TetraMag, LiDAR, IR, and visible light photogrammetry for a drone-based survey of the Glenn Dale inert minefield using the techniques developed in Chapters 2, 3, and 5. We begin by describing the field area, surveying methods, and steps taken to process the data. We then present TetraMag data, including anomaly maps generated by invariants of the finite-difference magnetic gradient tensor. After, we present FLIR data, including visible light and thermal photogrammetry. We then display the LiDAR data and consider the textural features of the produced point cloud. We combine these into multi-modal maps of the survey area. Finally, we discuss our approach's detection capabilities and limitations, including suggestions for future improvements.

6.1 Field Area Survey and Data Preparation

As described in Chapter 5, our field area consists of an inert minefield we constructed in Glenn Dale, MD (see Section 5.3). We buried 19 targets in a 17.5 m x 7.5 m portion of the field and included two additional control holes in the grid for a total of 21 known target locations. We flew a 0.5 m/s lawnmower grid over the field area with 0.3 m spacing at 0.5 m altitude to collect TetraMag and LiDAR data. We then flew the same grid at 10 m altitude to collect FLIR

thermal and visible light photogrammetry, as both cameras were out of focus at 0.5 m altitude. We then applied the data synchronization described in Chapter 5 to TetraMag data. We processed LiDAR data using GeoSLAM Draw, importing the high-precision GNSS measurements of the ground control points to georeference the 3D point cloud. We truncated the point cloud volume to a maximum of 0.3 m above the ground surface and filtered outliers from the dataset. We then exported the point cloud as a .laz file and imported it to MATLAB for visualization. We processed the FLIR thermal and visible light data using PIX4Dmapper. We separated the visible and thermal images from their combined RJPG file format. Then, we calibrated each set separately, calculating tie points using geometric similarity and location information from the camera’s attached GNSS puck. After initial processing, we merged the visible and thermal datasets for stitching. Since the thermal camera of the FLIR Duo has a much lower resolution than the visible light camera, we used the visible light data to calculate the 3D mesh geometry. We then projected the thermal texture on top, creating thermal and visible light orthomosaics.

6.2 TetraMag Data

Chapter 5, Figure 5.9 showed the field area before and after applying the noise model. We calculated the full magnetic gradient tensor (FDMGT) using the anomaly values and Equation 3.3. We then calculated the gradient tensor’s invariants for each timestep, as described in Section 3.3. Below, we show invariant maps for the eigenvalues, determinants (symmetric and antisymmetric parts), Frobenius norms (symmetric and antisymmetric, and the ratio of antisymmetric-to-symmetric), eigenvalue combinations 1 and 2, and NSS.

Figure 6.1 shows λ_1 (top), λ_2 (middle), and λ_3 (bottom) maps, with GNSS-measured target

positions overlaid. White "+" symbols indicate ferrous and possibly ferrous targets. White "o" symbols indicate non-ferrous targets or control holes. The λ_1 map shows anomalously low spots that align with targets A5-A7, B2, B4-B7, C1, and C3. λ_3 maps show anomalously high spots that align most strongly with targets A2, A4, A6, A7, B4, B6, and B7. Anomalously high and low spots for λ_2 maps align with the same target locations seen in maps λ_1 and λ_3 . All three eigenvalue maps present additional low and high anomalies not associated with the known targets.

Figure 6.2 shows $|G_S|$ (top), $\|G_S\|_F$ (middle), and $\|G_A\|$ (bottom) for the field area. Background signals are suppressed well for the symmetric determinant map, with targets A2-A4, A6, A7, B1, B2, B4-B7, C1, and C3 visible as spots with high or low anomalous spots compared with the background. Targets A2-A7, B1-B7, C1, C3, and C7 are localized as anomalously high spots for both $\|G_S\|_F$ and $\|G_A\|$ with amplitudes in the symmetric map nearly six times larger than the antisymmetric. Control holes and non-ferrous objects do not appear as anomalies in any maps.

As we showed in Chapter 3, the $\|G_A\|_F/\|G_S\|_F$ ratio is a potentially useful parameter for determining burial depth (see Figure 3.6). Figure 6.3 displays this ratio for the field area (top). The bottom panel shows the predicted target depth using the Frobenius norm ratio at the GNSS-marked locations. The red line is a 1:1 ($y = x$) reference. A strongly linear dependence ($r^2 = 0.99$) exists between the Frobenius norm ratio and the measured burial depth. Table 6.1 summarizes the measured burial depths, predicted burial depths using the Frobenius norm ratio, and the error as a percentage of the measured burial depth. While error percentages vary between 1 – 20%, all errors are less than 1 cm.

Figure 6.4 displays maps for λ_{C1} (top), λ_{C2} (middle), and NSS (bottom). Anomalously high spots localize targets A2-A4, A6, A7, B1, B2, B4-B7, C1, and C3 for λ_{C1} , while λ_{C2} localizes the same targets with a mix of both anomalously high and low spots. Targets A2-A7,




Figure 6.1: Eigenvalue anomaly maps from a drone-based survey of the inert minefield. A white "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A white "o" symbol indicates GNSS locations of non-ferrous targets and control holes.




Figure 6.2: Determinant and Frobenius norm anomaly maps from a drone-based survey of the inert minefield. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.




Figure 6.3: Frobenius norm ratio $\|G_A\|_F / \|G_S\|_F$ map of the inert minefield (top) and the corresponding predicted vs measured depth values (bottom). The red line in the bottom panel is a 1:1 reference line.

Target	Description	Measured Depth (cm)	Predicted Depth (cm)	Error (%)
A2	AT Mine	7	7.54	7.8
A3	Dummy Bomb	10	9.89	1.1
A4	IED Analog	13	12.71	2.2
A5	EM Mine	5	5.87	17.3
A6	Dummy Grenade	8	8.97	12.1
A7	Simulant AT mine	10	10.43	4.3
B1	M42 Submunition	10	10.91	9.1
B2	AT mine	10	9.51	4.9
B3	PDM AP mine	10	9.25	7.5
B4	Dummy Bomb	10	10.10	1.0
B5	Claymore AP	5	5.96	19.3
B6	Simulant AT mine	7	7.83	11.9
B7	Bouncing Betty AP mine	12	11.03	8.1
C1	AT mine	9	8.63	4.1
C3	Metal AT mine	10	9.34	6.6
C4	BLU-26 Bomb	10	10.70	7.0
C5	Artillery Shell Analog	5	5.13	2.7
C7	Simulant AT mine	8	7.88	1.6

Table 6.1: Detected ferrous target measured and predicted burial depths. Predicted depths are calculated using the Frobenius norm ratio $\|G_A\|_F/\|G_S\|_F$.

B2-B7, C1, C3, and C7 are localized by anomalously high spots for the NSS map. Control holes and non-ferrous objects do not appear as anomalies in any of the maps.

6.3 FLIR Data

Figure 6.5 shows the stitched visible light orthomosaic from the 10 m flight, where each pixel represents a single 4000x3000 photo that has been calibrated, zoomed, and rotated to match neighboring photo key features. Areas with anomalous features, defined as having sharp color contrasts with the rest of the photo or very symmetric features, are highlighted at insets on the orthomosaic. The top left inset shows an anomalous area where, once zoomed in, target B7 is visible. The green cross-hair position indicates the location with the most anomalous feature compared to its surroundings. The middle and right insets show circular patches where the grass is not growing, exposing the soil beneath. The four white squares in the photo are the survey area ground control points. Figure 6.6 shows the orthomosaic cropped around the survey area. Anomalous features are circled in black.

Figure 6.7 shows the stitched thermal orthomosaic for the inert minefield. Temperatures reported are relative, with 0 representing the coldest spots and 1 representing the hottest spots. White squares correspond to survey ground control points. Figure 6.8 shows the thermal orthomosaic cropped around the survey area, with anomalously high- or low-temperature features circled in white. Figure 6.9 shows the cropped survey area, where the anomalous features are isolated from the background signal.




Figure 6.4: Eigenvalue combinations and normalized source strength anomaly maps from a drone-based survey of the inert minefield. A black "+" symbol indicates GNSS locations of ferrous and possibly ferrous targets. A black "o" symbol indicates GNSS locations of non-ferrous targets and control holes.




Figure 6.5: Visible light orthomosaic of the larger field area. Survey markers appear as white boxes in the photo. Examples of detected anomalous areas are shown in the inset boxes. The top left inset shows target B7. The bottom center and right insets show areas where oval-shaped patches of soil differ from the surrounding vegetation (targets C3 and C1, respectively).




Figure 6.6: Visible light orthomosaic of the inert minefield. Anomalous features are in black circles.

6.4 LiDAR Data

Figure 6.10 shows the processed 3D LiDAR point cloud from oblique (a) and nadir (b) angles, and the color is scaled by LiDAR return intensity. Distance measurements are made with respect to the position of the GeoSLAM Zeb Horizon upon takeoff. The field area has a vertical gradient, visible in (a). Figure 6.11 shows the point cloud after being georeferenced using GNSS measurements from the four ground control points. Areas with anomalously low return intensity are circled in white.

6.5 Multi-Modal Anomaly Maps

In this section, we present anomaly maps using multiple modalities, using the LiDAR map as a base on which other data are draped. Figure 6.12 shows the visible light orthomosaic overlaid on the LiDAR intensity map. Anomalies in LiDAR alone are circled in white, while LiDAR and




Figure 6.7: Thermal orthomosaic of the field area. Temperature data are relative, where 0 is the coldest, and 1 is the hottest spot. White squares are ground control points.




Figure 6.8: Thermal orthomosaic of the inert minefield with anomalous features circled in white.




Figure 6.9: Thermal orthomosaic of the inert minefield with anomalous features isolated and circled in white.




Figure 6.10: Processed GeoSLAM Zeb Horizon LiDAR data. (a) shows an oblique view, and (b) a nadir view of the survey area. The coordinate system is referenced to the drone takeoff location, not the survey markers.




Figure 6.11: Georeferenced LiDAR map with anomalies circled in white.

visible light anomalies are circled in pink.

Figure 6.13 shows thermal and visible light orthomosaics draped on the LiDAR return intensity map. Anomalies only in thermal imagery are circled in yellow, while anomalies present in all three modalities are circled in pink.

Figure 6.14 shows thermal, visible light, and λ_2 invariant maps draped on LiDAR return intensity. Anomalies only in thermal imagery are circled in yellow, while anomalies only in the invariant map are circled in black. Anomalies that appear in multiple modalities are circled in pink. This same color designation continues throughout all the following multi-layer maps with magnetic gradient invariants. Figure 6.15 shows visible light, thermal, and $|G_S|$ invariant maps draped on LiDAR return intensity. Figure 6.16 shows visible light, thermal, and $\|G_F\|_F$ invariant maps draped on LiDAR return intensity, while Figure 6.16 shows the $\|G_F\|_F$ invariant. Figure 6.17 shows visible light, thermal, and λ_{C1} invariant maps draped on LiDAR return intensity, while Figure 6.18 shows λ_{C2} . Figure 6.19 shows the NSS invariant map, visible light, and thermal




Figure 6.12: Visible light orthomosaic draped on the LiDAR return intensity map. Anomalies present only in LiDAR are circled in white. Anomalies in both LiDAR and visible light are circled in pink.




Figure 6.13: Thermal and visible light orthomosaics draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow, and anomalies present in LiDAR, visible, and thermal orthomosaics are circled in pink.




Figure 6.14: Visible light, thermal, and λ_2 invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.

draped on LiDAR return intensity. Finally, all anomalies are layered in Figure 6.20. The radius of circles on this map represents the known location of each target, with 0.25 m tolerance. A positively identified target contains anomalies in at least one of the modalities within this range. Known target labels appear above the circled anomalies associated with each, with black circles corresponding to anomalies only seen in TetraMag, yellow only in thermal, and pink circles seen in multiple modalities. The locations of the control holes, A1 and C2, are in dashed black circles and are not detected as anomalies by any of the instruments. The dashed circles with question marks in the center correspond to the additional anomalies detected by TetraMag and/or thermal that are not associated with known targets.

6.6 Discussion

Here, we discuss the results above and their implications. We individually discuss each sensor's detection capabilities and limitations (FLIR, LiDAR, TetraMag) and then compare these




Figure 6.15: Visible light, thermal, and $|G_S|$ invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.




Figure 6.16: Visible light, thermal, and $\|G_S\|_F$ invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.




Figure 6.17: Visible light, thermal, and λ_{C1} invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.




Figure 6.18: Visible light, thermal, and λ_{C2} invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.




Figure 6.19: Visible light, thermal, and NSS invariant maps draped on the LiDAR return intensity map. Anomalies present only in thermal imagery are circled in yellow. Anomalies present only in the invariant map are circled in black. Anomalies present in multiple modalities are circled in pink.

to the stacked multi-modal anomaly maps.

6.6.1 Detection Capabilities and Limitations for Individual Sensors

Invariant maps produced from TetraMag data localize ferrous and low-ferrous targets very well in the field area, though not all equally. In particular, the eigenvalues (Figure 6.1) include many anomalous high and low spots not seen in other invariant maps and not coinciding with known target locations. $|G_S|$, seen in Figure 6.2 (top), resolves the individual targets better than λ_1 , λ_2 , or λ_3 . The detected targets are easier to distinguish from the background and are well located with respect to the known GNSS locations. The signal from A5 is weak but still present in $|G_S|$, which may suggest lower ferrous content. Frobenius norm invariants in Figure 6.2 (middle and bottom) localize the targets well, with $\|G_S\|_F$ producing higher amplitude anomalies that coincide with known ferrous target locations. The ratio of $\|G_A\|_F/\|G_S\|_F$ in Figure 6.3 reveals depth information that is consistent with the known burial depths within 2 cm, reinforcing the




Figure 6.20: Visible light, thermal, λ_{C2} , $|G_S|_F$, $|G_S|$, and NSS invariant maps draped on the LiDAR return intensity map.

importance of considering the antisymmetric part produced from the magnetic gradient tensor, especially when small-scale, small-amplitude signals are present. The eigenvalue combination and NSS maps in Figure 6.4 also localize the targets well, though signal amplitudes from λ_{C2} produce larger amplitude anomalies and suppress background signals slightly better than λ_{C1} , which may be tied to the orientation of each ferrous target's dipole. NSS anomaly maps show high anomalies over ferrous targets, consistent with the literature. By examining the invariants of the magnetic gradient tensor, anomalous signals from small targets are amplified, making their detection more likely than traditional methods. The main limitation of using TetraMag in isolation is that it is sensitive only to targets with at least some ferrous material. Not all landmine, UXO, or ERW casings are detectable with this method alone.

The visible light and thermal orthomosaics produced from the FLIR Duo Pro R are useful for geolocating and calibrating other sensors and detecting surface-laid munitions. The visible light orthomosaic seen in Figures 6.5 and 6.6 identify three anomalous regions, one of which was confirmed to be target B7 from a zoomed-in image, and the other two aligned with known buried target locations. The thermal orthomosaic localizes many targets and is especially useful as a confirmatory metric for other modalities such as TetraMag. Since the visible light camera has a much higher resolution, we used its 3D mesh as the base for the thermal map projection, with limited success initially. However, we achieved a reasonable result after changing some of the processing parameters. A slower-speed survey and/or tighter transect spacing is necessary for reliable results. Such high overlap requirements constitute a significant limitation of this technique when used in isolation.

The LiDAR return intensity maps are useful not only for draping the other sensor data on but also for their return intensity textural features. The LiDAR texture maps show anomalous

spots around known target locations A3, A5, A7, B2, B6, B7, C1, C3, C5, and C7. These textural features may arise from the ground disturbance from buried targets and the divots resulting from target burial. It is worth noting that there were two months between burial and surveying so that these divots may be a reliable detection metric for recently buried landmines, UXO, and ERW. The return intensity may also be low in burial spots where the overlying vegetation has not fully recovered from the target burial, as seen in the visible light photogrammetry. LiDAR texture alone does not recover the locations for targets A2, A4, A6, B1, B3, B4, B5, or C4, highlighting the limitation of LiDAR as a means of detection.

6.6.2 Detection Capabilities and Limitations of Multi-Modal Maps

Combining multiple modalities in a single anomaly map produced far more accurate and reliable results than the individual sensors could. For instance, combining LiDAR and visible light results in the positive detection of 10 targets. Including thermal data shows overlap of 5 targets, though the thermal data has distortion-related offsets (see Figures 6.12 and 6.13). Most known targets are detected once the TetraMag gradient tensor invariants are included. The most accurate results occur when combining $|G_S|$, $\|G_S\|_F$, λ_{C2} , or NSS with visible light, thermal, and LiDAR, as all expected targets are detected. Additional anomalies are seen in the invariant maps below target C7, which may be a piece of clutter associated with farming equipment. The presence of these targets in the thermal data helps verify the likelihood that those signals come from a real source and are not merely artifacts from the interpolation. Figure 6.20 shows the combination of visible light, thermal, λ_{C2} , $\|G_S\|_F$, $|G_S|$, and NSS invariant maps draped on the LiDAR return intensity map, with circles adjusted to center on the known GNSS target locations. Dashed cir-

cles A1 and C2 were not detected by any modality, demonstrating that back-filled holes without anthropogenic targets are not a source of erroneous detection. Target C6 was only detected with thermal photogrammetry (yellow circle), while targets A2, A4, and B3 were detected only with TetraMag invariant maps (black circles). Targets detected by multiple modalities are circled in pink, including A3, A5, A6, A7, B1, B2, B4, B5, B6, B7, C1, C3, and C7. Dashed circles with question marks in the center represent targets that appeared in multiple TetraMag invariants but were not part of the known set of buried objects. Since the thermal data also show anomalies here, we are confident they are real. While robust, this technique suffers from long computing time for visible and thermal photogrammetry, and careful survey planning is necessary to ensure each modality meets its individual requirements. Every modality included in the multi-layer map enabled the successful detection of all known targets in the survey area, including three additional anomalies of unknown origin. Since the additional anomalies appear in the TetraMag invariant maps, it can be reasonably assumed that, if real, they are ferrous.

6.7 Conclusions

This study demonstrates that multiple geophysical modalities enhance the detection of small-scale, small-amplitude anomalies such as landmines, UXO, and ERW. While LiDAR, visible light and thermal photogrammetry, and TetraMag magnetic gradient tensor invariants can detect some of the targets individually, the combination of all methods enabled the successful detection of all 19 buried targets, with a localization error tolerance of 3 cm. This technique is limited mainly by payload weight and computation power required to process some of the included modalities. The full instrument payload limits flying time to < 15 minutes per set of

batteries, which reduces the survey coverage capabilities. This can be improved by using lighter fluxgate magnetometer sensors with TetraMag. Processing the visible and thermal data is both time-consuming and computationally expensive. Processing data from a single survey took days, limiting how quickly one can produce the multi-modal maps. This could be improved using a camera or software other than PIX4Dmapper. Despite these limitations, the techniques developed here are robust. They can be applied to landmine, UXO, ERW detection and characterization efforts, and general critical zone geoscience. Other use cases for the payload package and data processing framework developed here include mapping the shallow critical zone, detecting non-lethal anthropogenic targets such as utilities, and hydrogeologic studies, among many others.

Chapter 7: Conclusions and Future Research

In this chapter, we summarize the main conclusions from this work, including lessons learned and directions for future research that build on our results. The key takeaway is that geophysical methods, data analysis techniques, and instrumentation enable the detection and characterization of small-scale, small-amplitude targets in the shallow critical zone, including landmines, UXO, and ERW. Chapter-specific conclusions are summarized below.

7.1 Chapter Conclusions

Chapter 1 introduced the global landmine, UXO, and ERW crisis as a humanitarian-focused use case of multimodal geophysics. However, the methods and instruments developed within are useful for general geophysical surveying, especially for the shallow critical zone, which is poorly understood. The techniques developed to address the issue must be geoscience-centered, as current methods that rely solely on engineering principles have been largely unsuccessful.

Chapter 2 described and discussed the geophysical sensors used in this project. Ground-based GPR is suboptimal for realistic minefield conditions due to data quality issues caused by needing to lift the instrument across difficult-to-navigate terrain. Utilizing commercial GPR on a drone is also not feasible due to power loss and side reflections caused by the choice of antenna. IR and Visible Light photogrammetry have been utilized for near-surface geophysics. They are

often used for landmine and ERW detection since their thermal properties are different than the natural environment. LiDAR is useful for the precision localization of targets but also provides textural feature information that can be leveraged to locate munitions or areas of interest with potential buried objects. TetraMag, which we developed in-house using MEDA FVM400 triaxial fluxgate sensors, successfully overcomes the limitations of single-sensor systems. By measuring the gradients across the triaxial sensors, one can construct the full finite-difference magnetic gradient tensor, which is rotationally invariant, insensitive to distant noise sources from space weather and cultural features, and sensitive to small-scale changes near the sensor. Our general approach is summarized in Figure 2.14. It includes the use of both geophysical and telemetry sensors to detect and characterize the small-scale, small-amplitude signals from landmines, UXO, and ERW.

Chapter 3 illustrated how the finite size of magnetic gradiometers compared to the length scales of magnetic anomalies due to small buried objects affects the symmetry and invariant properties of the finite-difference magnetic gradient tensor (FDMGT). This is in contrast to traditional assumptions of theoretical magnetic gradient tensors, which should be symmetric and whose invariants should not be affected by rotations. We showed that deviations from symmetry decrease with the square of the distance from the magnetic target (represented as a dipole). Still, they can exhibit complicated spatial patterns depending on the orientation of the magnetization of the target. Further, we showed how the rotational dependence of invariants persists even when using the symmetric component of the FDMGT.

We explored the utility and benefit of leveraging the symmetric and antisymmetric parts of the FDMGT for locating and characterizing small-scale, small-amplitude targets in the shallow critical zone. In particular, the Frobenius norm, NSS, and eigenvalue combinations λ_{C1} , λ_{C2}

enable high-resolution localization of small magnetic targets and characterization of their source parameters such as burial depth and magnetization orientation. Building on theoretical frameworks from potential fields literature, we built and tested a regular tetrahedral triaxial magnetic gradiometer called TetraMag and demonstrated its ability to estimate the FDMGT, then compared each derived invariant's ability to localize the buried target.

The methodology presented in Chapter 3 enables rapid processing and interpretation of field data. Since most of the errors between the IMGT and FDMGT are wrapped up in creating an antisymmetric part of the MGT, utilizing the symmetric part obviates the need for some of the more rigorous calibration efforts traditionally applied to magnetic survey data. Our approach helps lay the foundation for UAV-based detection of small targets. The relative stability of invariant patterns to rotations of TetraMag enables the instrument to be rigidly mounted on the vehicle body. This effectively shifts the correction task to the data analysis side rather than the field operational side, where gimbals and other equipment would traditionally be used always to keep the instrument level.

Chapter 4 demonstrates how machine learning techniques such as MLPs and CNNs are useful for detecting small-scale, small-amplitude targets and estimating their source parameters. When GPR data is available, the CNN from Section 4.2 may highlight radargrams of interest that likely contain UXO with 95% accuracy. Adding a sliding window to identify the detected UXO hyperbola's location automatically would make this method even more useful for rapidly identifying regions of interest. When a regression layer is added to the end of the CNN architecture, source parameters such as magnetic moment and burial depth may be estimated from FDMGT invariant maps with high accuracy and acceptable error margins. The CNN's ability to estimate both magnetic moment and burial depth highlights that important information is contained within

the FDMGT that is not necessarily captured by measuring just B_t , or the vector components B_x , B_y , and B_z , especially at these small scales. As a data processing tool, ML enables the fusion of multiple geophysical modalities when detecting and characterizing targets in the shallow critical zone, whether anthropogenic or geological.

Chapter 5 presented a method for reducing magnetic self-noise from a UAV by utilizing telemetry data for a rigidly-mounted triaxial magnetic gradiometer (TetraMag). It is standard for commercial drones to log yaw, pitch, roll, Ω , and battery status ($\vec{I}$), so no additional sensors are required to utilize the developed technique for any drone-mounted magnetic system. However, if desired, this does not preclude using extra, higher-resolution accelerometers and/or current sensors to measure these parameters. To measure $\vec{I}$ in a non-invasive way, one could attach split-core current sensors at each of the motor wire connections and log the data using either the flight controller or an external microcontroller fitted with a real-time clock. Data fidelity and spatial coherence improve dramatically for models A, B, and C compared to results using $B_{Exp}^{\rightarrow}$ only. Model C, which includes currents, motor speed, and their derivatives, ultimately outperforms the others. It is consistent with the conservation of energy equation (5.1) and the motor voltage balance equation (5.2), which govern DC motors. Measurable magnetic signals from the vibration of the drone may be suppressed effectively with a low-pass filter, further increasing data fidelity. This method may be improved by using higher sampling frequency magnetometers, enabling faster surveying velocity and wider coverage per set of drone batteries. For the Freefly Alta X, a higher sampling frequency for TetraMag would ensure no telemetry data are lost while downsampling. Since the Alta X has a vibration-damping ring between its body and our custom landing gear, it is worth investigating the extent to which yaw, pitch, and roll measured by the drone matches the movement of TetraMag. Integrating an additional accelerometer on the landing gear

would verify the accuracy of the telemetry measurements and inform future versions of the drone package. This method is not restricted to UAV-based magnetometry and may be applied to any vehicle that utilizes DC motors. Ground robotics systems such as the AeroVironment tEODor EVO [178], for example, which disarms and disposes of explosive ordnance, could be fitted with TetraMag and use these modeling methods to model and subtract out the vehicle's self-noise for UXO detection before removal.

Chapter 6 combined the results obtained in Chapters 1-5 and applied them to an inert mine-field we constructed in Glenn Dale, MD. This study demonstrated that multiple geophysical modalities enhance the detection of small-scale, small-amplitude anomalies such as landmines, UXO, and ERW. While LiDAR, visible light and thermal photogrammetry, and TetraMag magnetic gradient tensor invariants are able to detect some of the targets individually, the combination of all methods enabled the successful detection of all 19 buried targets, with a localization error tolerance of 3 cm. This technique is limited mainly by payload weight and computation power required to process some of the included modalities. The full instrument payload limits flying time to < 15 minutes per set of batteries, which reduces the survey coverage capabilities. This can be improved by using lighter fluxgate magnetometer sensors with TetraMag. Processing the visible and thermal data is both time-consuming and computationally expensive. Processing data from a single survey took days, limiting how quickly one can produce the multi-modal maps. This could be improved by using a different camera or software other than PIX4Dmapper. Despite these limitations, the techniques developed here are robust. They can be applied to landmine, UXO, and ERW detection and characterization efforts and to general critical zone geoscience. Other use cases for the payload package and data processing framework developed here include mapping the shallow critical zone, detecting non-lethal anthropogenic targets such as utilities,

and hydrogeologic studies, among many others.

7.2 Lessons Learned

Here, we describe the challenges faced during this project and the lessons we learned.

7.2.1 Sourcing and Operating High-Payload Drones

One of the biggest challenges we faced during this project was finding a drone with sufficient payload capacity to suit our requirements and learning to operate it without the support of experienced pilots and/or aerospace engineers. If given the opportunity to re-do this project, I would ensure that the Freefly Alta X was acquired within the first year of the project and that sufficient training was arranged and provided before attempting any activities with the vehicle. The learning curve for drone piloting was steep and would have only improved if given more time. Not having access to the drone until the end of year 4 of the project hindered noise modeling efforts, though the results we achieved are sufficient for our purposes.

7.2.2 The Importance of Appropriate Sampling Frequency

While the relatively low cost of the MEDA 400FVM enabled us to build a four-sensor gradiometer, the limited sampling frequency presented challenges throughout this work. The maximum sampling frequency we achieved while in continuous data streaming mode was 4 Hz, dramatically reducing the drone velocity required to have sufficient data coverage and fidelity of our various field areas. One of the lessons we learned was to prioritize a higher sampling frequency fluxgate magnetometer to increase survey speed and more fully leverage the higher

sampling frequency telemetry sensor data.

7.2.3 The Importance of Magnetometer Calibration

Talk about needing to subtract off the mean from all magnetometers when conducting corner surveys and how this is tied to magnetometer 2 having an offset from the other magnetometers tied to being out of calibration. While de-meaning works in case calibration is not possible, it is best practice to ensure fluxgate magnetometers are calibrated annually, especially if they have been saturated by being in highly magnetic areas. Our magnetometers have been operated in various settings, including inside a highly magnetized indoor laboratory. This may have saturated the magnetometers and necessitates calibration before future surveys.

7.2.4 Limiting Factors for Drone-Based Geophysics

Payload weight and battery capacity are the most limiting factors for any drone-based work, especially when combining both commercial and bespoke sensors as a combined payload. With the Alta X fully equipped, we could fly for an average of 10 – 13 minutes for each battery set, which decreased surveying efficiency. To achieve longer flight times that are more appropriate for geophysical surveying, the payload weight must be decreased. One way to do this would be to re-design TetraMag to either use different, more lightweight sensors or to prioritize the analog data rather than digital, which could reduce the weight per sensor by half. A potential downside to this approach would be an increase in necessary post-processing steps to ensure proper filtering and calibration. Still, the increase in flying time alone makes it inconsequential.

7.2.5 The Challenge of Relying on Open-Source Tools

We utilized open-source software and hardware wherever possible during this project. This was the optimal choice for some applications, such as collecting and organizing TetraMag data on the open-source Raspberry Pi. For other applications, however, such as processing LiDAR and FLIR data, open-source software was both difficult to find and not well supported, easily quadrupling the time it would take to process those data using proprietary software packages such as GeoSLAM Draw (for LiDAR) and PIX4Dmapper (for FLIR).

7.3 Future Research

Here, we present ideas for follow-up research based on the results of this work.

7.3.1 UAV-Optimized Triaxial Fluxgate Magnetometers

The design and fabrication of a more payload-friendly vector magnetic gradiometer with a higher sampling frequency would be a useful follow-on project to enable longer flying times, faster flying speeds, and better data coverage. Altering the MEDA 400FVM sensors to achieve this is possible. Still, it may require the assistance of an electrical engineering student to work on the hardware side, while a geophysics student can build on our codes for the software side. Optimizing such a sensor suite for geoscience purposes requires the participation and active research from a geophysicist, who understands what good magnetic data look like and how to process them, which is key to making progress.

7.3.2 Utilizing Flight Controller for Instrument Control

An additional future project could involve utilizing the drone's flight controller to sequence and control geophysical instruments. Due to time constraints, we could not fully integrate each sensor into the Alta X's Black Cube, but such a task is feasible, especially for TetraMag and FLIR. The Black Cube has additional input and output spaces that could control the timing of FLIR photos and trigger TetraMag's data collection and storage functions. This would obviate the need for an additional real-time clock and streamline data processing steps since the data would be synchronized automatically. This would also pave the way to near real-time data processing aboard the drone.

7.3.3 Quantifying Optimal Survey Conditions and Correcting FLIR Data with an Albedometer

The quality of thermal data depends on several factors, including the reflectivity of the survey area. A follow-on project could integrate a lightweight albedometer into the drone package to better quantify this and find optimal surveying conditions. By comparing the data quality with albedo measurements, one could determine the best time of day to conduct a FLIR survey and potentially correct for atmospheric interference in the IR images (or identify them as anomalous to streamline data processing better).

7.3.4 Integrating Mid-Wavelength FLIR Camera for Gas Leak Detection

The FLIR Duo Pro R used in this study is considered a long-wavelength IR device (8 – 15 μm). This is appropriate for detecting larger-scale thermal differences between landmines and the surrounding soil. Adding a mid-wavelength (3 – 8 μm) thermal camera would allow one to detect gas leaks. It would be a useful tool for a methane mapping project at finer scales than the Cessna-mounted detection system can provide.

7.3.5 Drone-Based Software-Defined GPR

One of the sub-projects we worked on that ended up being too time-intensive for the scope of this thesis is the development of a software-defined GPR compatible with low-altitude flying drones. While commercial GPR units are ill-suited for drone work, the technique is ripe for research and development for geophysical surveying. Based on the work of [84, 179–182], we tested the capabilities of NanoVNA, Analog Devices ADALM-PLUTO, and AMD RFSoc 4x2 software defined radio hardware as potential signal generation devices for a drone-based GPR. We found that the RFSoc 4x2 had the most robust capabilities and would be lightweight and low power enough to be used on the Freefly Alta X. We performed initial antenna design tasks and found a pair of Vivaldi antennas to be the most promising. The follow-on project would involve picking up where we left off to write the PYNQ GPR signal generation and processing codes and designing/fabricating the appropriate instrument housing to be attached to the drone. A student could then test the imaging capabilities using known reflectors buried in the relatively homogeneous sandbox, utilizing the rolling scaffold setup we constructed at the Paint Branch Turfgrass Facility to explore the data fidelity when using pulsed-echo, frequency sweeps, and a "comb"

of simultaneous frequencies. Since the RFSoc 4x2 has 4 RF ADCs and 2 RF DACs, one could also explore the extent to which polarization data (either circularly or linearly polarized) enables shape and dielectric permittivity characterization of detected targets. The drone implementation of SDR GPR on the Alta X may be as simple as extending lateral booms at the vibration-damping disk and mounting the Vivaldi antennae on opposite sides of the drone. However, shielding may be necessary if it is found that the other payload instruments produce noise or side clutter in the measured radar signals.

7.3.6 Shallow Critical Zone Moisture Mapping with Drone-Based SDR GPR and FLIR

Once the SDR GPR is made, a series of projects could combine the GPR with the previously integrated LiDAR, FLIR, and TetraMag. One project would be to use lower frequency signals generated by the SDR GPR alongside FLIR to map shallow subsurface soil conductivity. Since radar waves attenuate with increasing soil moisture, the data quality and produced radargrams could be leveraged to characterize the soil conductivity across a field area. The FLIR could be utilized to highlight pockets of higher soil moisture, as they would have a different thermal response than the surrounding soil. This has potential tie-ins for investigating the impact of road salt use near local watersheds, for example.

7.3.7 Detection and Mapping of Environmental Contaminants

A project that could utilize drone-based FLIR and TetraMag data would be detecting and mapping environmental contaminants such as steam condensate leaks. For steam leaks, one

could fly over areas with suspected leaks to first map out the location of ferrous pipes carrying the steam, as they would have different magnetic and thermal properties than the surrounding soil. Once the pipes have been approximately located, the next step would be to conduct drone FLIR surveys before dawn, at dawn, and after dawn, as one can expect fresh leak locations to have a continuous influx of warmer material over time, and areas farther from the active leak to exhibit thermal decay over time.

7.3.8 Mapping of Oxidation Fronts for Magnetic Minerals

Magnetic minerals, such as magnetite, maghemite, hematite, and goethite, form in the soil in response to complex pedogenic processes that are sensitive to physical, chemical, and biological conditions and are useful for paleoclimate reconstruction and climate change studies [183]. TetraMag could map the ferrous oxide fronts in a field area if flown at low altitudes. Total field magnetic gradiometry has been used to study soil creep and pedoturbation associated with the weathering of ferrous basalt dikes [184]; invariants of the finite-difference magnetic gradient tensor TetraMag produces may provide more information to understand the evolution of these weathering features. These data may also be useful for tracking the geomorphological history of an area [185].

7.3.9 Flight Path Optimization using FDMGT Invariants

As demonstrated in Chapter 3, the ratio of the antisymmetric-to-symmetric parts of the magnetic gradient tensor's Frobenius Norm is proportional to the distance from the sensor to the magnetic source. A follow-on project would leverage this information to update a flight plan,

effectively seeking out more ferrous material rather than relying on a standard grid and post-processing to update the plan. This could be especially useful when mapping large anomalies such as abandoned wellheads from the oil industry.

Appendix A: Exploring Principal Component Analysis, Convolutional Neural Networks, and Scattering Transforms in a Rock Image Classifier

Advances in instrumentation, storage technology, computational power, and compression for storage management have enabled geosciences to collect large datasets. As the size and collection rate of geoscientific data increase, so does the need to automate some classification tasks that help us make inferences from the data. Despite advances in machine learning (ML) performance and reliability, methods are often portrayed as “black boxes” without statistical or physical meaning. This leads to hesitation in incorporating ML as a data processing tool in scientific workflows. Traditional dimensionality reduction techniques such as principal component analysis (PCA) remain ubiquitous in the geosciences. Still, ML-based methods like convolutional neural networks (CNNs) and scattering transforms (STs) promise high accuracy and meaningful feature extraction for classification tasks. This work aims to (1) evaluate PCA against the ML-based methods of CNNs and ST summary coefficients as inputs to a multilayer perceptron (MLP) network for classifying rocks, (2) gain insight into how each technique discriminates between rock type classes, and (3) explore the extent to which differences in the ability to distinguish rock types arise when using grayscale versus color images. We designed a simple rock classifier MLP using two fully-connected, hidden layers and an output layer. The input dataset consisted of 29,998 28×28 pixel hand-sample images of andesite, gneiss, marble, quartzite, rhyolite, and schist, in

grayscale and 3-channel color versions. We performed PCA, CNN, and ST calculations with two different orientations on each rock image, then used 80% of the outputs to train the MLP. After training, the remaining 20% were used to evaluate the performance of each network. To visualize and understand the classification performance of each technique, we used t-Distributed Stochastic Neighbor Embedding (t-SNE) on the PCA, CNN, and ST summary coefficient outputs. The ST-based rock classifier outperforms PCA by at least 20% accuracy in both grayscale and color and outperforms the CNN by 1 – 4% accuracy in grayscale and color. Negligible differences in accuracy are observed when using 2 versus 8 ST orientations.

A.1 Machine Learning in Geosciences

As the ability to rapidly collect large, high-quality datasets increases, so does the need to automate some classification tasks that help us make inferences from the data. Classification problems in the geosciences include but are not limited to discriminating rock types from a large sample set, characterizing fracture system topologies, classifying volcanic material to infer eruption mechanisms, and categorizing diverse seismic events [186]. Despite advances in ML performance and capability, impressive results are often reported without insight into the methods that created them. This leads to hesitation in incorporating ML as a standard data processing tool [186].

We evaluate the traditional dimensionality reduction technique of Principal Component Analysis (PCA) against the ML-based methods of convolutional neural networks and scattering transforms when all are used as inputs to the same multilayer perceptron network for rock classification. Since image-based data are not always 3-channel color, an additional goal is to explore

how differences in the ability to distinguish rock types arise when using grayscale versus color images.

A.1.1 Multi-Layer Perceptron (MLP)

As discussed in Chapter 4, multilayer perceptron (MLP) is a machine learning technique that has its origin in Rosenblatt’s 1957 report, “The Perceptron, A Perceiving and Recognition Automation,” which aimed to design an analog of the human brain for data analysis [145]. MLPs can be used unaccompanied for classification problems but are often incorporated at the back end of other machine learning techniques in modern applications [149]. Some examples of MLPs in the geosciences include predicting swell potential of clay-rich soils [187], distinguishing between volcanic earthquakes and false events [188], evaluating geochemical parameters in organic data [189], and detecting faults from 3D seismic data [190].

A.1.2 Principal Component Analysis (PCA)

Principal Component Analysis (PCA) is a multivariate statistics technique that captures essential data patterns by finding “lines and planes of closest fit to systems of points in space” and projects the important information in a set of orthogonal variables called principal components (eigenvectors of the covariance matrix), which are used to find patterns of similarity [191]. The goals of PCA include simplification, data reduction, modeling, outlier detection, variable selection, classification, and prediction [192]. PCA is non-parametric and linear in how it reduces dimensionality, meaning that mapping from higher-dimensional space to lower-dimensional space happens via matrix multiplication, so it cannot capture higher-order dependencies that may exist

in the input data nor interpret complex polynomial relationships between features [193]. Examples of PCA used in the geosciences include de-noising seismic data [194], simplifying geochemical datasets for exploration and mining geology [195], separating hydrochemical facies and anomalous fluoride content from GIS and tectonic data [196], separating crossing events and diffracted waves from reflected waves in seismic data [197], and identifying changes in surface features from satellite data [198].

A.1.3 Convolutional Neural Networks (CNNs)

Convolutional Neural Networks (CNNs) are one of the most commonly used types of artificial neural networks, designed to replicate how humans process visual data under various physical transformations [150]. They are conceptually similar to MLPs but are translationally and rotationally invariant [151]. Additionally, CNNs introduce nonlinearity and pooling layers that adjust, limit, and summarize the generated outputs, allowing for robust and efficient pattern recognition, especially in image-based datasets [152]. Functionally, CNNs learn and dynamically adapt spatial hierarchies of features by passing an optimized feature extractor over smaller subsections of the data, summarizing the dominant features, and associating them with ground truth labels that are used to update weights and biases of the classification network. Examples of CNNs in the geosciences include automating P-wave arrival picking to generate focal mechanisms [199], classifying volcanic ash particles based on their shape to infer eruption mechanisms [200], predicting spatial groundwater potential for hydrological maps [201], calibrating subsurface flow system models [202], and producing high-resolution geological maps from unmanned aerial vehicle (UAV) data [203].

A.1.4 Scattering Transform (ST)

The scattering transform (ST) [204] is a relatively new technique of pattern recognition that computes transnationally-invariant coefficients that retain high-frequency information and are stable under diffeomorphisms via wavelet modulus decompositions and low-pass averaging filters [205]. Representations produced in the scattering transform method have a similar structure to CNNs [206]. Unlike CNNs, which use the Fourier power spectrum to represent stationary processes at lower order [207], the ST is Lipschitz-continuous to random deformation [204]. By neglecting the phase, spectral correlations that may have otherwise been hidden from higher-order moments are captured by the ST, making it an effective method of texture recognition and discrimination [206].

While a CNN will produce filter banks with pooling functions and nonlinearities that are learned while training by the chosen gradient descent method, the ST will produce a bank of wavelet filters and pointwise nonlinearities via stability, invariance, and informative requirements. The ST will also produce a feature representation at every level, allowing for direct interpretation at any layer in the network; a CNN will only produce a feature representation at the final layer. Tunable hyperparameters of the ST implemented in the MATLAB-based ScatNet software include scale (J), which controls the timescale of the lowpass filter, orientations (L), which controls the number of wavelet orientations, and order (M), which controls the maximal order of the transform [208].

Uses of ST in the geosciences include categorizing seismic structures from seismic images [209], predicting stratigraphic units from well logs [210], and detecting and clustering seismic events from continuous seismic data [211].

A.2 Data Organization and Methods

The dataset we used is publicly available from Kaggle and consists of 29,998 3-channel hand-sample images of andesite, gneiss, marble, quartzite, rhyolite, and schist, each of which is 28x28 pixels in size [212]. We split the set into 80% training and 20% testing, with equal amounts of rock sample images from each class. We separated images by folder hierarchy, each representing one of the 6 rock types. We then loaded the images into MATLAB in grayscale and 3-channel color formats. The 3-channel color images were separated by red, green, and blue channels, and the 4 versions of each image (grayscale, red channel, green channel, and blue channel) were saved as separate vectors.

A.2.1 PCA Methods

For grayscale images, we performed PCA using MATLAB native commands, which normalize the input prior to calculating principal components. 33 total components were calculated per image, resulting in a 23998×33 training matrix and a 6000×33 testing matrix, both of which were then fed in to the MLP.

We treated each color in the 3-channel color images separately, and PCA was performed in MATLAB. 33 total components were calculated for each channel (red, green, and blue), then combined for 99 principal components per image. The resulting training matrix was 23998×99 , and the testing matrix was 6000×99 , each of which was then fed to the MLP.




Figure A.1: Convolutional Neural Network (CNN) architecture. Each of the 23998 rock images is fed in and traverses through two convolution layers, with max pooling functions between them, then flattened at the output layer. The flattened outputs are then fed into the MLP for classification.

A.2.2 CNN Methods

We constructed the CNN in Python using Tensorflow, SciKitLearn, and Keras libraries [213]. The architecture of the CNN consists of 2 convolution layers with max pooling and ReLu (rectified linear unit) activation functions between each convolution layer (see Figure A.1). Layer 1 contains 48 filters with 5×5 kernel size and 14×14 outputs, and layer 2 contains 33 filters with 3×3 kernel size, resulting in 7×7 outputs. After passing through the convolution layers, the output is flattened before being fed to the MLP. Using the original 3-channel color image dataset, only one channel was used for grayscale image runs, while all three channels were used for color image runs.

A.2.3 ST Methods

We used Scatnet, a MATLAB-based software, to calculate scattering transforms for both grayscale and 3-channel color versions of the rock sample images with two different configurations: (a) $J = 4$, $L = 2$, $M = 2$, and (b) $J = 4$, $L = 8$, $M = 2$, where J represents timescale for the

lowpass filter, L represents number of orientations, and M represents maximal order.

For grayscale images with $J = 4$, $L = 2$, and $M = 2$, a total of 33 ST coefficients were produced, each with 7×7 dimension. We then calculated the absolute mean value of each of the 7×7 ST coefficients to produce 33 single summary coefficients for each image. The resulting training matrix was 23998×33 , and the testing matrix was 6000×33 , each fed to the MLP.

For grayscale images with $J = 4$, $L = 8$, and $M = 2$, a total of 417 ST coefficients were produced, each with a 7×7 dimension. We then calculated the absolute mean value of each of the 7×7 ST images to produce 417 single summary coefficients for each image. The resulting training matrix was 23998×417 , and the testing matrix was 6000×417 , each of which was then fed to the MLP.

For 3-channel color images with $J = 4$, $L = 2$, and $M = 2$, each channel was treated separately (red, green, and blue), and a total of 33 ST coefficients were produced per channel, each with 7×7 dimension. We calculated the absolute mean value of each of the 7×7 coefficients to produce 33 single-value ST summary coefficients, then combined them into a 1×99 vector of ST summary coefficients for each image. The resulting training matrix was 23998×99 , and the testing matrix was 6000×99 , each fed to the MLP.

For 3-channel color images with $J = 4$, $L = 8$, and $M = 2$, each channel was treated separately, and a total of 417 ST coefficients were produced per channel, each with 7×7 dimension. We calculated the absolute mean value of each of the 7×7 coefficients to produce 417 single-value ST summary coefficients, then combined them into a 1×1251 vector for each image. The resulting training matrix was 23998×1251 , and the testing matrix was 6000×1251 , each fed to the MLP.

A.2.4 t-Distributed Stochastic Neighbor Embedding (t-SNE) for Visualization

To gain insight into how PCA and ST methods organize data in higher dimensions after performing their respective operations, we used t-Distributed Stochastic Neighbor Embedding (t-SNE) as a visualization aid. t-SNE is an unsupervised, nonlinear technique that minimizes the divergence between the distribution of pairwise similarities in the inputs and the distribution of measurements of pairwise similarities in corresponding lower-dimensional points [214]. The power of incorporating t-SNE is in its ability to balance attention between local and global patterns in the data via its perplexity hyperparameter, mapping a faithful representation to lower-dimensional space.

A.2.5 MLP Methods

We used each of the competing methods (PCA, CNN, and ST) the input to an MLP consisting of two fully connected layers and an output layer (see Figure A.2).

Using single channel (grayscale) input, the first fully connected layer contained 33 nodes for PCA, CNN, and ST with $J = 4$, $L = 2$, and $M = 2$. For the ST with $J = 4$, $L = 8$, and $M = 2$, the first fully connected layer contained 417 nodes to accommodate the additional coefficients calculated using 8 orientations rather than 2. Using 3-channel (color) images, the first fully connected layer contained 99 nodes, with 33 nodes for each channel, for PCA, CNN, and ST with $J = 4$, $L = 2$, and $M = 2$. For the ST with $J = 4$, $L = 8$, and $M = 2$, the first fully connected layer contained 1251 nodes to accommodate the additional coefficients calculated with 8 orientations rather than 2. All first fully connected layers used ReLu as the activation function.

The second fully connected layer contained 48 nodes and a ReLu activation function for




Figure A.2: Multilayer Perceptron (MLP) architecture. Outputs from either PCA, CNN, or ST are used as input to the MLP, where they traverse through two fully-connected layers containing 33 and 48 nodes, respectively, before reaching the output layer, containing 6 nodes corresponding to the 6 rock types (andesite, gneiss, marble, quartzite, rhyolite, and schist). Positive edges between nodes are shown in yellow, and negative edges are shown in orange

single- and 3-channel versions of PCA, CNN, and ST.

The output layer contained 6 nodes, representing the 6 potential rock types: andesite, gneiss, marble, quartzite, rhyolite, and schist. Since this is a multiclass problem, a softmax activation function was applied at the output layer.

We used sparse categorical cross-entropy as the loss function for the MLP, with Adam optimization and accuracy evaluation metrics. The MLP was allowed to train for 100 epochs using a batch size of 60. No additional hyperparameters were tuned across the PCA, CNN, and ST during or after training.

A.3 Results

In this section, we present the results from each of the ML methods described above.

A.3.1 PCA Results

Beginning with grayscale rock images, 33 principal components were calculated and used as input to the MLP. After training, the highest accuracy on the testing dataset was 72.2% (see Figure A.3a). While the training accuracy converged to approximately 85% by Epoch 100, the validation accuracy converged around Epoch 20 and remained relatively stable throughout the following 80 Epochs. Loss decreased throughout all 100 training epochs but increased during validation after epoch 20 and continued to increase in epochs thereafter (see Figure A.3b). Figure A.4 shows the confusion matrix for the model when fed input images unseen during the training and validation stages. The model misclassifies rhyolite as marble 3.03% of the time, quartzite as andesite 2.08% of the time, schist as gneiss 2.83% of the time, and schist as quartzite 2.35% of




Figure A.3: Accuracy and Loss of PCA-based Rock Classifier (Grayscale Images) (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.

the time.

Using 3-channel color rock images, 33 principal components were calculated for each channel, combined to make 99 total principal components, and used as input to the MLP. After training, the highest accuracy was 87.4% (see Figure A.5a). The training accuracy converged to approximately 99% by Epoch 60, but the validation accuracy was less stable and decreased from an average of 88.5% in the first 20 epochs to 87.5% by epoch 60. Loss decreased over the 100 training epochs but increased almost linearly during validation (see Figure A.5b). The confusion matrix shows the model misclassifies schist as gneiss 2.38% of the time, rhyolite as andesite 2.18% of the time, gneiss as schist 1.87% of the time, and andesite as quartzite 1.97% of




Figure A.4: Confusion matrix of PCA-based Rock Classifier (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.




Figure A.5: Accuracy and loss for PCA-based Rock Classifier (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.

the time (see Figure A.6).

A.3.2 CNN Results

Using grayscale rock images, the inputs underwent two rounds of convolution and max pooling before being flattened and fed into the MLP. After training, the highest accuracy was 90.5% (see Figure A.7a). The training accuracy increased throughout its 100 epochs, converging to 99% by epoch 95. Validation accuracy converged to 90% by epoch 40 and remained relatively stable throughout the following 60 epochs, with some jitter around epoch 65 (see Figure A.7a). Training loss consistently decreased throughout its 100 epochs, but validation loss slightly in-




Figure A.6: Confusion matrix for PCA-based Rock Classifier (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.




Figure A.7: Accuracy and loss for CNN-based Rock Classifier (Grayscale Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.

creased after epoch 60, with a notable spike in loss around epoch 65 (see Figure A.7b). Figure A.8 shows the confusion matrix when the model was fed images unseen in training and validation. The model misclassifies quartzite as andesite 1.83% of the time, andesite as quartzite 1.65% of the time, and schist as gneiss 1.22% of the time (see Figure A.8).

Using 3-channel (color) rock images, the inputs underwent two rounds of convolution and max pooling before being flattened and fed into the MLP. After training, the highest accuracy reached using images previously unseen by the model was 98% (see Figure A.9a). Both training and validation accuracies increased throughout each of the 100 epochs, with some jitter between epochs 15 and 80 (see Figure 9A.9a). Loss decreased for training and validation stages throughout




Figure A.8: Confusion matrix for CNN-based Rock Classifier (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.




Figure A.9: Accuracy and loss for CNN-based Rock Classifier (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.

100 epochs, with jitter at the same epochs as the accuracy plots (see Figure A.9b). The confusion matrix shows the model misclassifies quartzite as andesite 0.47% of the time, schist as gneiss 0.43% of the time, andesite as quartzite 0.33% of the time, and rhyolite as andesite 0.23% of the time, and andesite as rhyolite 0.23% of the time (see Figure A.10).

A.3.3 ST Results

Using grayscale rock images as input to scatnet with $J = 4$, $L = 2$, $M = 2$, 33 scattering transform coefficient images were produced, each with 7×7 pixel dimensions. The average absolute value of each 7×7 image was calculated to produce 33 summary coefficients for each




Figure A.10: Confusion matrix for CNN-based Rock Classifier (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.

image, which were then fed as input to the MLP. After training, the highest accuracy reached on rock images unseen during training and validation was 93.3% (see Figure A.11a). Accuracy increased across all 100 epochs for both the training and validation stages, with negligible jitter for validation accuracy occurring between epochs 10 and 45 (see Figure A.11a). Loss consistently decreased across 100 epochs for training and validation stages of model creation; small deviations in validation loss occur near epochs 41 and 48 (see Figure A.11b). Figure A.12 shows the confusion matrix for the previously unseen rock images. The model misclassifies quartzite as andesite 1.70% of the time, schist as gneiss 1.10% of the time, andesite as quartzite 0.92% of the time, andesite as rhyolite 0.72% of the time, and rhyolite as andesite 0.57% of the time (see Figure A.12).

Using grayscale images as input to scatnet with $J = 4$, $L = 8$, $M = 2$, 417 scattering transform coefficient images were produced, each with 7×7 pixel dimensions. The average absolute value of each 7×7 image was calculated to produce 417 summary coefficients, which were then fed as input to the MLP. After training, the highest accuracy on rock images unseen during training and validation was 93.8% (see Figure 13a). Accuracy increased for training and validation stages across 100 epochs, with the validation accuracy converging to 98% by epoch 78 (see Figure A.13a). Loss decreased across all 100 epochs for training but slightly increased after epoch 80 for the validation stage (see Figure A.13b). The confusion matrix shows the model misclassifies quartzite as andesite 2.02% of the time, andesite as quartzite 0.97% of the time, schist as gneiss 0.67% of the time, andesite as rhyolite 0.58% of the time, and rhyolite as andesite 0.48% of the time (see Figure A.14).

Using 3-channel color rock images as input to ScatNet with $J = 4$, $L = 2$, and $M = 2$, 33 scattering transform coefficients were produced separately for the color channels, each with




Figure A.11: Accuracy and loss for ST-based Rock Classifier, $L = 2$ (Grayscale Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.




Figure A.12: Confusion matrix for ST-based Rock Classifier, $L = 2$ (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.




Figure A.13: Accuracy and loss for ST-based Rock Classifier, $L = 8$ (Grayscale Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.




Figure A.14: Confusion matrix for ST-based Rock Classifier, L = 8 (Grayscale Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.

7×7 pixel dimensions. The average absolute value of the 7×7 images was calculated to produce 33 summary coefficients for each channel, combined in a single vector for 99 total summary coefficients, and fed as inputs to the MLP. After training, the highest accuracy reached using rock images unseen during training and validation was 98.1% (see Figure A.15a). Accuracy increased across 100 epochs for training and validation, with minimal jitter for the training stage (see Figure A.15a). Loss decreased for training and validation across 100 epochs, with some jitter for the validation stage after epoch 80 (see Figure A.15b). Figure A.16 shows the confusion matrix when using previously unseen rock images and shows the model misclassifies andesite as quartzite 0.57% of the time, rhyolite as andesite 0.42% of the time, schist as gneiss 0.33% of the time, quartzite as andesite 0.17% of the time, and andesite as rhyolite 0.13% of the time.

Using 3-channel color rock images as input to ScatNet with $J = 4$, $L = 8$, and $M = 2$, 417 scattering transform coefficients were produced separately for the color channels, each with 7×7 pixel dimensions. The average absolute value of the 7×7 images was calculated to produce 417 summary coefficients for each channel, combined in a single vector of 1251 summary coefficients, and fed as inputs to the MLP. After training, the highest accuracy reached using rock images unseen during training and validation was 98% (see Figure A.17). The confusion matrix produced with previously unseen rock images as inputs shows the model misclassifies schist as gneiss 0.83% of the time, quartzite as andesite 0.47% of the time, rhyolite as andesite 0.38% of the time, andesite as quartzite 0.20% of the time, and andesite as rhyolite 0.10% of the time (see Figure A.18).




Figure A.15: Accuracy and loss for ST-based Rock Classifier, $L = 2$ (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.




Figure A.16: Confusion matrix for ST-based Rock Classifier, L = 2 (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.




Figure A.17: Accuracy and loss for ST-based Rock Classifier, $L = 8$ (Color Images). (a) shows the accuracy for training (red) and testing (blue) as a function of epoch number. (b) shows the loss of training (red) and testing (blue) as a function of epoch number.




Figure A.18: Confusion matrix for ST-based Rock Classifier, L = 8 (Color Images). Input images were not used during the training or validation phases. True labels correspond to ground truth, and the model generates predicted labels after running through the trained MLP.

A.4 Discussion

This section discusses the results of each technique for grayscale and color images.

A.4.1 PCA: Grayscale vs Color

Figures [A.19](#) and [A.20](#) show t-SNE plots of calculated principal components for grayscale and color images.

A.4.1.1 PCA with Grayscale Rock Images

Using grayscale images results in lower MLP accuracy than the color image version for the PCA-based rock classifier (see Figure [A.3](#)). From Figure [A.19](#), which shows the t-SNE projection for grayscale images, limited distinguishable clustering occurs using 33 principal components. Some separation is visible for marble and gneiss, which would suggest an MLP may be able to discriminate the two classes more easily, and the confusion matrix supports this claim, as there are no instances where gneiss is misclassified as marble nor marble misclassified as gneiss (see Figure [A.4](#)). The t-SNE projection for rhyolite and marble show considerable intermingling, suggesting the MLP may confuse the two. The confusion matrix supports this claim, where the highest instances of misclassification occur when marble is misidentified as rhyolite (3.90%), and rhyolite is misidentified as marble (3.03%) (see Figure [A.4](#)). Quartzite and schist show similar intermixing and are misidentified frequently by the MLP, with quartzite being misclassified as schist 2.30% of the time and schist as quartzite 2.35% of the time (see Figures [A.4](#), [A.19](#)). Schist is more frequently misidentified as gneiss than gneiss as schist (2.83% vs 1.42%), which is consistent with the t-SNE plot, where schist appears spread out along the entire manifold shape, with limited




Figure A.19: t-SNE projection of PCA scores (Grayscale Images).

clustering (see Figure A.19).

A.4.1.2 PCA with 3-Channel Color Rock Images

Using 3-channel rock images results in 15.2% higher MLP accuracy than grayscale images for the PCA-based rock classifier (see Figures A.3 and A.5). Figure A.20 shows the t-SNE representations for each channel and the projection for the combined principal components. There is limited clustering or visible separation between each rock class when viewing the channels separately. Still, substantial clustering is present for the combined 3-channel t-SNE representation (see Figure A.20, lower panel). Marble has a distinct grouping with well-defined boundaries compared to the t-SNE representations in grayscale, suggesting the MLP would misidentify that class less frequently. The confusion matrix for the 3-channel PCA rock classifier confirms this, as marble has the fewest instances of misclassification across all rock types (see Figure A.6).

Quartzite and rhyolite appear to have distinct boundaries despite being neighboring clusters, suggesting the MLP would misclassify them less frequently (see Figure A.20). The confusion matrix confirms this, as rhyolite is misclassified as quartzite 0.43% of the time, and quartzite as rhyolite 0.02% of the time (see Figure A.6). Schist and gneiss, however, show considerable intermixing in the t-SNE representation, suggesting the MLP might misclassify the two. The confusion matrix supports this suggestion, as schist is misclassified as gneiss 2.38% of the time and gneiss as schist 1.87% of the time. Obscured from the static view of Figure A.20, rhyolite and andesite are substantially intermixed, suggesting they may be difficult to distinguish, which is confirmed by the confusion matrix, where rhyolite is misidentified as andesite 2.18% of the time. Surprisingly, andesite is misclassified as rhyolite only 0.40% of the time, suggesting additional distinguishing features may be present for andesite samples (see Figure A.6).

A.4.2 CNN: Grayscale vs Color

Figures A.21 and A.22 show t-SNE plots of outputs from the CNN for grayscale and color images after the second convolution layer and the last layer.

A.4.2.1 CNN with Grayscale Rock Images

Using grayscale images as input, the CNN-based rock classifier accuracy reached 90.5%, outperforming the PCA-based versions for both grayscale and 3-channel inputs (see Figures A.4, A.6, and A.8). Figure A.21a, the t-SNE representation for the CNN output following convolution layer 2, contains no discernable clustering for the six rock type classes. This is unsurprising, as CNNs do not produce feature representations until their final layer. However, the t-SNE projec-




Figure A.20: t-SNE projections of PCA scores for 3-channel input images. PCA was performed on red, green, and blue channels separately, then combined to create the 3-channel principal components. t-SNE representations for the separate channels are shown in the upper portion, and the representation for the combined principal components is shown below

tion of the output at the final layer shows very little separation of the outputs. It does not elucidate how the CNN can separate classes from one another with such high accuracy (see Figure A.21b). While layer-by-layer feature maps can be produced for an image by running it individually through a pre-trained network, visualizing the operations the CNN performs during training would help develop an intuition for what features are important from input data to discriminate between different classes.

A.4.2.2 CNN with 3-Channel Color Rock Images

Using 3-channel color versions of the input rock images, the CNN-based classifier accuracy reached 98.0%, outperforming the grayscale CNN by 7.5% and the 3-channel PCA classifier by 10.5% (see Figures A.6, A.8, and A.10). Figure A.22a, the t-SNE representation for the CNN output following convolution layer 2, shows a different manifold shape than the grayscale version but still contains no discernable clustering of the six rock-type classes. Similarly, the t-SNE representation before the final layer shows a different manifold shape than the second convolution layer but no clustering of individual classes with one another (see Figure A.22b). It is unclear what physical meaning, if any, can be gained from observing how the t-SNE projection changes for the outputs of each CNN layer. However, it is interesting to note that the number of group clusters matches the number of nodes for the final convolution layer in Figure A.22b.

A.4.3 ST: Grayscale vs Color

Figures A.23-A.26 show t-SNE plots of ST summary coefficients calculated for grayscale and color images, for $J = 4$, $L = 2$, $M = 2$, as well as $J = 4$, $L = 8$, $M = 2$.




Figure A.21: t-SNE projections for CNN extracted features after the second convolution layer (a) and before the last layer (b) for grayscale input images.




Figure A.22: t-SNE projections for CNN extracted features after the second convolution layer (a) and before the last layer (b) for 3-channel color input images.

A.4.3.1 ST with Grayscale Rock Images

Using grayscale images and summary coefficients calculated with 4 timescales, 2 orientations, and 3 orders ($J = 4$, $L = 2$, $M = 2$), the ST-based rock classifier reached 93.3% accuracy, outperforming the grayscale CNN classifier by 2.8% and the grayscale PCA classifier by 21.1% (see Figures A.4, A.8, and A.12). Substantial clustering for most of the 6 classes is visible in the t-SNE plot, suggesting the MLP may perform better than the grayscale PCA classifier (see Figure A.23). Andesite and gneiss show distinct boundaries and distance from one another in the t-SNE representation, suggesting the MLP may have an easier time distinguishing between the two, and the confusion matrix supports this, as there are no instances where andesite is misclassified as gneiss, nor gneiss as andesite (see Figure A.12). Schist and marble are similarly separated from one another, and the confusion matrix supports that the MLP should have an easier time discriminating between the classes, as there are no instances of schist being misclassified as marble, nor marble as schist (see Figure A.12). There is considerable intermingling between andesite and quartzite clusters, suggesting the MLP may struggle to distinguish one from the other; this is supported by the confusion matrix, where quartzite is misclassified as andesite 1.70% of the time, and andesite is misclassified as quartzite 0.92% of the time. Schist and gneiss show some clustering and intermixing, and the schist cluster is more spread out than gneiss, suggesting not only that the MLP may confuse the two but that it may misclassify more often when schist is the ground truth label, rather than gneiss (see Figure A.23). The confusion matrix supports these claims, as schist is misclassified as gneiss 1.10% of the time, but gneiss is misclassified as schist only 0.58% of the time (see Figure A.12). Andesite and rhyolite clusters intermix at some positions and are misclassified as one another 0.72% of the time, for andesite as rhyolite, and 0.57%




Figure A.23: t-SNE plot for summary ST coefficients with $J = 4$, $L = 2$, $M = 2$ (Grayscale Images).

of the time (for rhyolite as andesite).

Using grayscale images and summary coefficients calculated with 4 timescales, 8 orientations, and 3 orders ($J = 4$, $L = 8$, $M = 2$), the ST-based rock classifier reached 93.5% accuracy, outperforming the grayscale CNN classifier by 3%, the grayscale PCA classifier by 21.3%, and the $L = 2$ ST classifier by 0.2% (see Figures A.6, A.10, A.14, and A.16). Similar to the $L = 2$ case, substantial clustering is visible in the t-SNE representation for most of the 6 rock type classes, though the $L = 8$ case has a different manifold shape than $L = 2$ (see Figures A.23 and A.24). Quartzite has multiple smaller clusters spread throughout the t-SNE plot, and one of the quartzite clusters mixes with the andesite cluster near the center, suggesting the MLP may confuse the two (see Figure A.24). The confusion matrix supports this, as quartzite is misclassified as andesite 2.02% of the time, and andesite is misclassified as quartzite 0.97% of the time (see Figure A.16). Quartzite being misidentified as andesite more frequently than andesite as quartzite is consistent




Figure A.24: t-SNE plot of summary ST coefficients with $J = 4$, $L = 8$, $M = 2$ (Grayscale Images).

with the observation that fewer andesite clusters are present in the t-SNE plot, and the clusters are closer to one another. Schist and gneiss clusters show some intermixing, suggesting the MLP may confuse the two; the confusion matrix shows schist is misclassified 0.67% of the time, and gneiss is misclassified as schist 0.53% of the time (see Figure A.16). Obscured from the static view of Figure A.24, rhyolite and andesite clusters also exhibit substantial intermixing and are misidentified as one another 0.48% and 0.58% of the time, respectively.

A.4.3.2 ST with 3-Channel Color Rock Images

Using 3-channel color images and summary coefficients calculated with 4 timescales, 2 orientations, and 3 orders ($J = 4$, $L = 2$, $M = 2$), the ST-based rock classifier reached 98.1% accuracy, outperforming the 3-channel CNN classifier by 0.1%, the 3-channel PCA classifier by 21.2%, and the 3-channel $L = 2$ ST classifier by 4.8% (see Figures A.6, A.10, A.16, and

A.18). Figure A.25 shows t-SNE plots for the separate channel summary coefficients and the combined color coefficients. Though clustering is present for t-SNE plots of each color channel, the combined summary coefficients result in a t-SNE plot with considerable clustering and larger separation between clusters than the grayscale ST and color PCA versions (see Figures A.20, A.23, and A.25). There is a substantial distance between rhyolite and marble clusters, suggesting the MLP may distinguish the two classes more easily; this is supported by the confusion matrix, as there are no instances where rhyolite and marble are misclassified as one another (see Figure A.16). Gneiss and rhyolite are similarly separated in Figure A.25, and there are no misclassification instances between the two classes in the confusion matrix (see Figure A.16). Andesite has multiple, smaller clusters throughout the t-SNE representation, some of which are mixed with quartzite clusters, suggesting the MLP may misclassify the pair; indeed, andesite is misclassified as quartzite 0.57% of the time, and quartzite is misclassified as andesite 0.17% of the time. The gneiss cluster shows good separation from other classes, except near the center, where there is intermixing with schist, whose cluster is more spread out, suggesting not only that the classes may be confused by the MLP but that misclassification is likelier when the ground truth is schist. This is confirmed by the confusion matrix, where schist is misclassified as gneiss 0.33% of the time, but gneiss is misclassified as schist only 0.13% of the time (see Figure A.16).

Using 3-channel color images and summary coefficients calculated with 4 timescales, 8 orientations, and 3 orders ($J = 4$, $L = 8$, $M = 2$), the ST-based rock classifier reached 98.0% accuracy, matching the performance of the 3-channel CNN classifier, outperforming the 3-channel PCA classifier by 21.1%, and underperforming by 0.1% when compared to the 3-channel $L = 2$ ST classifier (see Figures A.6, A.10, A.16, and A.18). Similar to the $L = 2$ case, the individual channel t-SNE plots have distinguishable clustering present for the six rock-type classes. However,




Figure A.25: t-SNE projections of ST summary coefficients for 3-channel input images, $J = 4$, $L = 2$, $M = 2$. STs and summary coefficients were calculated on red, green, and blue channels separately and then combined to create the 3-channel ST summary coefficients. t-SNE representations for the separate channels are shown in the upper portion, and the representation for the combined coefficients is shown below.

the combined summary coefficient t-SNE plot has more well-defined clusters with larger separations (see Figure A.26). A substantial distance between marble and andesite clusters suggests that the MLP may have more success distinguishing between the two (see Figure A.26). This is confirmed by the confusion matrix, where there are no instances where marble and andesite are misclassified as one another (see Figure A.18). Gneiss and andesite are similarly separated in t-SNE space, and there are also no instances of misclassification between the two. There is apparent intermixing between schist and quartzite clusters in the static view of Figure A.26. However, a more distinct separation is visible when the view is rotated, suggesting the MLP may be able to separate the boundary between the two more easily. The confusion matrix confirms this, as there are no instances of misclassification between schist and quartzite. Schist is more spread out than other clusters, and there are schist mini-clusters inside the gneiss cluster; this suggests the MLP may confuse the two, especially when schist is the ground truth. The confusion matrix confirms this, as schist is misclassified as gneiss 0.83% of the time, and gneiss as schist 0.00% of the time (see Figure A.18). The rhyolite cluster has distinguishable boundaries, with some mixing present in the andesite cluster, suggesting the two may be confused by the MLP; the confusion matrix shows rhyolite misidentified as andesite 0.38% of the time and andesite as rhyolite 0.10% of the time.

A.4.3.3 ST Orientations ($L = 2$ vs $L = 8$)

The accuracy of the ST-based rock classifier with 2 orientations is very similar to those with 8 orientations for both grayscale and 3-channel color rock images. Computationally, however, the two orientations perform very differently: on average, the $L = 8$ calculations took 75% longer to




Figure A.26: t-SNE projections of ST summary coefficients for 3-channel input images, $J = 4$, $L = 8$, $M = 2$. STs and summary coefficients were calculated on red, green, and blue channels separately and then combined to create the 3-channel ST summary coefficients. t-SNE representations for the separate channels are shown in the upper portion, and the representation for the combined coefficients is shown below.

run for both grayscale and color versions.

A.5 Conclusions

PCA produces consistently lower accuracies than CNN and ST for grayscale and 3-channel color rock images. This suggests that texture may be an important factor in discriminating rock samples from one another. Since textures are often the result of non-Gaussian processes, characterizing them by their spectral densities alone, as PCA does, may not be sufficient and is likely one of the reasons the ST technique outperforms PCA and, to a lesser extent, CNN. To compare how each technique performed with the same MLP design, no hyperparameters were tuned beyond the initial network design, so it is possible that ill-performing techniques like PCA could achieve higher accuracies after tuning hyperparameters.

The use of grayscale images results in lower accuracies for all techniques. Building on the assertion that texture features are critical for distinguishing between rock classes, color features provide sufficient information to help distinguish between rock types indistinguishable in grayscale. The supplementary computation time required to use three channels versus a single channel pays off in higher accuracies overall.

For the scattering transform technique, using 2 orientations versus 8 orientations results in little difference in accuracy, but 8 orientations requires 75% more computation time. While this suggests that fewer orientations are sufficient for this classification problem and dataset, it would be useful to test on different datasets and explore if the observation generalizes well, as it would have implications for storage handling on larger-scale projects.

Visualization of intermediate layer outputs with t-SNE provides useful information about

how PCA and ST extract meaningful, distinguishing features from each image. It may be used to estimate how well each input class may be accurately identified or which classes are likely to result in misidentification. Since CNNs do not output spatial mapping at each layer, t-SNE provides no additional insight or visualization for understanding how the convolution layers summarize extracted features, making it less accessible to a broader audience. While the CNN and ST performed similarly in accuracy, ST lends itself better to directly understanding the relationship between the summary coefficients and their original images via t-SNE. At the same time, the CNN's inner workings remain elusive.

Incorporation of t-SNE in ML network design and workflow practices may be useful for estimating potential performance. It would allow end-users to adjust their network designs before initiating any training. One could, for example, do initial PCA, ST, and other statistical summary calculations on a dataset, visualize the results with t-SNE, and then use that information to decide between potential classification or prediction architectures, saving time by eliminating the techniques whose t-SNE plots have poorly-defined clusters.

Appendix B: Standard Operating Procedures and Checklists

This appendix provides standard operating procedures for TetraMag, GeoSLAM Zeb Horizon, FLIR Duo Pro, as well as flight checklists for use when operating Debbie.

B.1 TetraMag SOP

1. Laptop/Peripherals Setup

- (a) Connect the Wi-Fi router ethernet cable to the laptop.
- (b) Connect the Wi-Fi router USB power cable to the laptop.
- (c) Check that Wi-Fi is broadcasting

SSID: HDDODY

Password: doandti23

2. Instrument Setup

- (a) Check that all RJ45 connectors are connected and secured at each sensor head and each dackie.
- (b) Check that all dackie 3mm-to-serial connectors are plugged in at the dackie and serial-to-USB hub.

- (c) Check that the serial-to-USB hub is plugged into the Raspberry Pi.
- (d) Switch dackies on using toggle.
- (e) Ensure battery levels are full (note: $\square\square\square\square x\square\square\square\square$ in the lower left of LCD screen indicates battery levels for both 9V are full $\square x\square$ in the lower left of the dackie LCD indicates battery levels for both 9V are $<25\%$).
- (f) Check that each dackie LCD displays data for X, Y, and Z components
- (g) Press ALT + Remote buttons on each dackie to put it into text output mode.
- (h) Check that the bottom right on dackie LCD displays “TEXT”.
- (i) Power on Raspberry Pi using the power bank (it should connect to the WiFi network automatically after booting).

3. Data Collection

- (a) On the laptop, open the command prompt and type:

```
ssh hddody@192.168.8.2
```

- (b) Once prompted for the password, type:

```
doandt123
```

- (c) To start data collection, type:

```
cd serial
```

```
sudo python3 TetraMag_MagOnly.py
```

Note: the command prompt screen should display something like:

```
Mag1: numbers, numbers, numbers
```

```
Mag2: numbers, numbers, numbers
```

```
Mag3: numbers, numbers, numbers
```

```
Mag4: numbers, numbers, numbers
```

If you do NOT see numbers displayed for all Mags, proceed to the troubleshooting section.

(d) To stop data collection and close the data file, type:

```
ctl+c
```

4. Data Offloading

- (a) Connect a monitor to the Raspberry Pi (note: "Serial" folder will contain all collected data)
- (b) Filename has a starting timestamp in UTC.
- (c) Copy/paste to the storage device of choice.

5. Troubleshooting

- (a) Dackie not turning on
 - i. Check batteries and replace them if needed.

(b) Magnetometer data not streaming to laptop control prompt screen

- i. Ensure all the cables are connected and secured.
- ii. Ensure all the dackies are in “TEXT” mode.
- iii. Turn relevant dackies off/on again, put in text mode and try again.

(c) SSH connection not initializing

- i. Check if the Wi-Fi router is broadcasting.
- ii. See if SSID shows up as an available network on a phone or other device.
- iii. Plug the touchscreen monitor into Raspberry Pi and check network connectivity.

If all else fails:

- i. Connect touchscreen monitor to Raspberry Pi
- ii. Open a terminal using Raspberry Pi GUI and type:

```
cd serial  
  
sudo TetraMag_MagOnly.py  
  
{ctl+c to stop}
```

B.2 FLIR Duo Pro R

1. FLIR Setup

- (a) Ensure GPS puck is plugged in and unobstructed.
- (b) Plug LiPO battery into FLIR Duo Pro R.
- (c) Open the ”FLIR UAS” app (preferably on an Apple device).



Figure B.1: FLIR Duo Pro Operation Buttons.

- (d) Press the Bluetooth button on the Duo Pro.
 - i. Ensure both SD cards are full on the app home screen.
 - ii. Press "record" on the app. This will initiate 1 Hz data collection.

B.3 GeoSLAM Zeb Horizon

Step 1	Connecting the hardware: Ensure battery is charged and connected to data logger. Connect the ZEB Horizon main cable to the socket on the side of ZEB Horizon scanner head. Connect the other end of the ZEB Horizon cable to the SCANNER (orange) socket on the ZEB Horizon data logger. Position the ZEB Horizon scanner on a flat stationary surface. Ensure the scan head is free to rotate. Turn on the ZEB Horizon data logger by pressing the on/off button.	Data Logger LED		Scan head LED array
		STATUS	DATA	
Step 2	Data logger booting up.		-	-
Step 3	Connecting to scan head.		-	
Step 4	Standby mode		-	
Step 5	Initiate new scan Long press the function button on either the scan head or the data logger until the STATUS LED and scan LED array light turns solid red to initiate a new scan			
Step 6	Initialisation mode The scan head must remain stationary for a period of 15 seconds. If the scan head is disturbed during initialisation the system will revert to standby mode (Step 4)		-	
Step 7	Scanning mode After the 15 second initialisation, the STATUS LED and scan head LED array will light green and the scan head will start rotating. Pick up the scanner and conduct the scan.		-	
Step 8	End scan - Data formatting To end the scan, long press the function button on either the scanner head or the data logger until the scan head stops rotating. The DATA LED will light orange while the scan data is converted to the required output format			
Step 9	Standby mode From stand-by either Initiate a new scan (Step 5), Download the data from the latest scans (section 8), or Shut down (section 8)		-	

Figure B.2: GeoSLAM Zeb Horizon SOP

PREFLIGHT CHECKLIST

Date:

Pilot in control (PIC):

0.	Weather	Check https://www.wunderground.com/dashboard/pws/KMDBOWIE82	
1.	Payload	SECURED	
1a	Landing Gear	SECURED	
2.	Mr. Squishy	VERIFY SECURE	
3.	Propellers	CHECK CONDITIONS, VERIFY TIGHT	
4.	Propeller Hubs	VERIFY SECURE	
5.	Motors	CHECK CONDITION	
6.	Radio Controller	ON, VERIFY TX BATTERY	
7.	Radio Controller Model	SELECT	
8.	Aircraft Placement	AWAY FROM PEOPLE AND OBSTACLES	
9.	Battery Pack Voltage	VERIFY ABOVE 48V	
10	Battery Packs	SECURE	
11.	Battery Leads	CHECK CONDITION AND CONNECT BATTERY #1	
12	Aircraft	KEEP STATIONARY	
13	Flight Controller	ALLOW TO INITIALIZE	
14	Blade Dampers	VERY 8 PRESENT AND IN GOOD CONDITION	
15	Orientation Lights	VERIFY SOLID COLORED, DIM	
16	Receiver	VERIFY BOUND	
17	ALTA QGroundControl	CONNECT, VERIFY NO WARNINGS	
18	Compass Calibration	CALIBRATE IF REQUIRED	
19	Radio Control Range Check	AS REQUIRED	
20	Battery	CONNECT BATTERY #2	

BEFORE TAKEOFF

Date:

Pilot in control (PIC):

1.	Prop Area	CLEAR	
2.	Mode Switch	POSITION	
3.	Home Switch	UP/OFF	
4.	Radio Controller	VERIFY CORRECT MODEL	
5.	Telemetry	CHECK OPERATION	
6.	GPS Signal	LOCKED (GREEN LIGHT)	
7.	ALTA X	CHECK ORIENTATION AND ARM	
8.	Orientation Lights	VERIFY USER-DEFINED COLOR	
9.	Motors	START, AND VERIFY OPERATION	
10.	Flight Controls	VERIFY CORRECT	
11.	Throttle	VERIFY MINIMUM	

POSTFLIGHT CHECKLIST

Date:

Pilot in control (PIC):

1.	Mode Switch	POSITION	
2.	Home Switch	OFF	
3.	ALTA X	LAND	
4.	Motors	DISARM AND STOP	
5.	Orientation Lights	VERIFY LIGHTS DIM	
6.	ALTA X QGroundControl	CHECK FOR WARNINGS	
7.	Batteries	DISCONNECT AND REMOVE	
8.	Radio Controller Power	TURN OFF	
9.	Aircraft Condition	INSPECT	
10.	Motor and Prop Condition	INSPECT	
11.	Battery Condition	INSPECT	
12.	Landing Gear	INSPECT	

Bibliography

- [1] Hany Kasban, O Zahran, Sayed M Elaraby, M El-Kordy, and Fathi E Abd El-Samie. A comparative study of landmine detection techniques. *Sensing and Imaging: An International Journal*, 11:89–112, 2010.
- [2] GM Turner, JL Rasson, CV Reeves, and M Kono. Observation and measurement techniques. *Treatise in Geophysics, Geomagnetism*, 5:91–135, 2015.
- [3] David Miles, Ian Mann, Andy Kale, David Milling, Barry Narod, John Bennest, David Barona, and Martyn Unsworth. The effect of construction material on the thermal gain dependence of a fluxgate magnetometer sensor. *Geoscientific Instrumentation, Methods and Data Systems Discussions*, pages 1–34, 2017.
- [4] Murat Kılınç, Osman Torunoğlu, and Hasan Koçer. The effect of viewing angle on detection of landmines from thermal time series images using active thermography. *International Journal of Energy Applications and Technologies*, 4(1):1–6, 2017.
- [5] Katherine James, Gert Riemersma, and Pedro Pacheco. How uav lidar imaging can locate and map minefield features: Cuito cuanavale, angola. *The Journal of Conventional Weapons Destruction*, 27(2):7, 2023.
- [6] Jasper Baur, Gabriel Steinberg, John Frucci, and Anthony Brinkley. An accessible seeded field for humanitarian mine action research. *The Journal of Conventional Weapons Destruction*, 27(3):2.
- [7] Inc. Freefly, 2023. url: <https://freefly.gitbook.io/freefly-public/products/alta-x/untitled-3/performance-specs>.
- [8] Raquel Salazar Moreno, Oscar Zerecero Salazar, et al. An artificial neural network model to analyze maize price behavior in mexico. *Applied Mathematics*, 9(05):473, 2018.
- [9] Ms Syeda Shabnam Hasan. *DETECTION OF EARLY ALZHEIMER’S DISEASE APPLYING MULTI-LAYER NEURAL NETWORKS*. PhD thesis, Ahsanullah University of Science and Technology, 2019.
- [10] Bubryur Kim, N. Yuvaraj, K. R. Sri Preethaa, and R. Arun Pandian. Surface crack detection using deep learning with shallow cnn architecture for enhanced computation. *Neural Computing and Applications*, 33:9289–9305, 1 2020.

- [11] M. Zucchetti, M. Khoder, I. Makki, R. Younes, C. Francis, and T. Bianchi. Landmines: Crisis, legacy, international and local action. *2017 1st International Conference on Landmine: Detection, Clearance and Legislations, LDCL 2017*, pages 22–27, 2017.
- [12] Kevin Bonsor. How landmines work, 2001. <https://science.howstuffworks.com/landmine.htm>.
- [13] Cluster Munition Coalition. Cluster munition monitor 2023. *ICBL-CMC*, 2023.
- [14] Robert Keeley. Understanding landmines and mine action. pages 2–7, 2003.
- [15] ICBL. Improvised mines: Casualties, 2019. <http://www.the-monitor.org/media/3073847/Improvised-Mines-Info-11'19.pdf>.
- [16] Mike Croll. *Landmines in War and Peace: From Their Origin to Present Day*. Pen and Sword Military; Illustrated edition, 2009.
- [17] ICBL. International campaign to ban landmines: A history of landmines, 2009. <http://www.icbl.org/en-gb/problem/a-history-of-landmines.aspx>.
- [18] Morris Panner. Landmines: A deadly legacy. *Fordham International Law Journal*, 308:1110, 1993.
- [19] Lesley Wexler. The international deployment of shame, second-best responses, and norm entrepreneurship: The campaign to ban landmines and the landmine ban treaty. *Arizona Journal of International and Comparative Law*, 20:561–606, 2003.
- [20] Henrique Siniciato Terra Garbino. The impact of landmines and explosive remnants of war on food security: the lebanese case. *Journal of Conventional Weapons Destruction*, 23(2):21–26, 2019.
- [21] Kathy Kowalenko. Saving lives, one land mine at a time. *IEEE The Institute*, 28:6, 2004.
- [22] Juan Carlos Ruan and Jill Macheme. Landmines in the sand : The falkland islands. *Journal of Conventional Weapons Destruction*, 5:43–45, 2001.
- [23] South Atlantic News Agency. Falkland islands land mines, how they are cleared and blown up — mercopress, 2019. <https://en.mercopress.com/2019/05/17/falkland-islands-land-mines-how-they-are-cleared-and-blown-up>.
- [24] Reuters. Colombia coca eradication hampered by landmines and protests, police say — reuters, 2020. <https://www.reuters.com/article/us-colombia-drugs/colombia-coca-eradication-hampered-by-landmines-and-protests-police-say-idUSKCN20L2NA?feedType=RSS>.
- [25] David J. Daniels. A review of gpr for landmine detection. *Sensing and Imaging*, 7:90–123, 2006.
- [26] Ursin Hoffman and Pascal Rapillard. Do no harm in mine action: Why the environment matters. *Journal of Conventional Weapons Destruction*, 19:13, 2015.

- [27] Y Li and M Nabighian. *Tools and Techniques : Magnetic Methods of Exploration – Principles and Algorithms*. 2015.
- [28] Gary D. Egbert, Markus Eisel, O. Sierra Boyd, and H. Frank Morrison. Dc trains and pc3s: Source effects in mid-latitude geomagnetic transfer functions. *Geophysical Research Letters*, 27:25–28, 1 2000.
- [29] K. N. Kappler, H. Frank Morrison, and G. D. Egbert. Long-term monitoring of ulf electromagnetic fields at parkfield, california. *Journal of Geophysical Research*, 115, 2010.
- [30] Andreas Nilsson, Neil Suttie, Monika Korte, Richard Holme, and Mimi Hill. Persistent westward drift of the geomagnetic field at the core-mantle boundary linked to recurrent high-latitude weak/reverse flux patches. *Geophysical Journal International*, 222:1423–1432, 2020.
- [31] H Robert Burger. Introduction to applied geophysics: Exploring the shallow subsurface pdf.
- [32] C P Gooneratne, S C Mukhopahyay, and G Sen Gupta. A review of sensing technologies for landmine detection : Unmanned vehicle based approach. *2nd International Conference on Autonomous Robots and Agents*, pages 401–407, 2004.
- [33] Jan M. H. Hendrickx, Nicole Alkov, Sung ho Hong, Remke L. Van Dam, Jan Kleissl, Heather Shannon, John Meason, Brian Borchers, and Russell S. Harmon. New mexico tech landmine, uxo, ied detection sensor test facility: measurements in real field soils. *Detection and Remediation Technologies for Mines and Minelike Targets XI*, 6217:62170V, 2006.
- [34] H.H. Nelson, Thomas W. Altshuler, E.M. Roesn, J.R. McDonald, B. Barrow, and N. Khadr. Magnetic modeling of uxo and uxo-like targets and comparison with signatures measured by mtads. *Proc. UXO Forum*, pages 282–291, 1998.
- [35] Jacqueline MacDonald, J. R. Lockwood, John McFee, Thomas Altshuler, and Thomas Broach. *Alternatives for Landmine Detection*. 1 2003.
- [36] L. Nuzzo, G. Alli, R. Guidi, N. Cortesi, A. Sarri, and G. Manacorda. A new densely-sampled ground penetrating radar array for landmine detection. *Proceedings of the 15th International Conference on Ground Penetrating Radar, GPR 2014*, pages 969–974, 2014.
- [37] Lee Sun Yoo, Jung Han Lee, Yong Kuk Lee, Seom Kyu Jung, and Yosoon Choi. Application of a drone magnetometer system to military mine detection in the demilitarized zone. *Sensors*, 21, 2021.
- [38] Yan Zhang, Leslie Collins, Haitao Yu, Carl E. Baum, and Lawrence Carin. Sensing of unexploded ordnance with magnetometer and induction data: Theory and signal processing. *IEEE Transactions on Geoscience and Remote Sensing*, 41:1005–1015, 2003.
- [39] NIST. Proton gyromagnetic ratio, 2021. <https://physics.nist.gov/cgi-bin/cuu/Value?gammap>.

- [40] H. Liu, H. Dong, Z. Liu, J. Ge, W. Luo, Z. Yuan, J. Zhu, and H. Zhang. A comprehensive study on the weak magnetic sensor character of different geometries for proton precession magnetometer. *Journal of Instrumentation*, 13:1–16, 2018.
- [41] Fritz Primdahl. The fluxgate magnetometer. *Journal of Physics E: Scientific Instruments*, 12(4):241, 1979.
- [42] Jack M Janicke. Second harmonic fluxgate sensors and magnetometers. *Sensor Review*, 18(4):225–229, 1998.
- [43] B. Gavazzi, P. Le Maire, J. Mercier de Lépinay, P. Calou, and M. Munsch. Fluxgate three-component magnetometers for cost-effective ground, uav and airborne magnetic surveys for industrial and academic geoscience applications and comparison with current industrial standards through case studies. *Geomechanics for Energy and the Environment*, 20:100117, 2019.
- [44] Songrui Wei, Xiaoqi Liao, Han Zhang, Jianhua Pang, and Yan Zhou. Recent progress of fluxgate magnetic sensors: Basic research and application. *Sensors*, 21:1–18, 2021.
- [45] Slawomir Tumanski. Induction coil sensors - a review. *Measurement Science and Technology*, 18, 2007.
- [46] Ground Penetrating Radar. Ground penetrating radar: theory and applications. *Jol ISBN*, pages 978–0, 2009.
- [47] a P Annan. Gpr-history,trends and future developments. *Subsurface Sensing Technologies and Applications*, 3:253–270, 2002.
- [48] Christian Wolff. Frequency-modulated continuous-wave radar (fmcw radar), 2021. [https://www.radartutorial.eu/02.basics/Frequency Modulated Continuous Wave Radar.en.html](https://www.radartutorial.eu/02.basics/Frequency%20Modulated%20Continuous%20Wave%20Radar.en.html).
- [49] Sabikun Nahar. *Design and Implementation of a Stepped Frequency Continuous Wave Radar System for Biomedical Applications*. PhD thesis, University of Tennessee, Knoxville, 2018.
- [50] Ahmet Serdar. Ground penetrating radar, 2015. <https://slideplayer.com/slide/4032282/>.
- [51] Hansel J. Martinez, Sthefany Alvarez, and Manuel A. Yarleque. Assessing the performance of three type of uwb antennas for fmcw gpr imaging. *Proceedings of the 2018 20th International Conference on Electromagnetics in Advanced Applications, ICEAA 2018*, pages 620–622, 2018.
- [52] Claudio Bruschini and Bertrand Gros. A survey of research on sensor technology for landmine detection. *The Journal of Humanitarian Demining*, 2:1–25, 1998.
- [53] Barry J. Allred. A gpr agricultural drainage pipe detection case study: Effects of antenna orientation relative to drainage pipe directional trend. *Journal of Environmental and Engineering Geophysics*, 18:55–69, 2013.

- [54] Steven C. Fisher, Robert R. Stewart, and Harry M. Jol. Ground penetrating radar (gpr) data enhancement using seismic techniques. *Journal of Environmental and Engineering Geophysics*, 1:89–157, 1996.
- [55] Cedric Schmelzbach and Emanuel Huber. Efficient deconvolution of ground-penetrating radar data. *IEEE Transactions on Geoscience and Remote Sensing*, 53:5209–5217, 2015.
- [56] Rasaan Bello. Literature review on landmines and detection methods. *Frontiers in Science*, 3:27–42, 2013.
- [57] Wim de Jong and John Schavemaker. Using polarization features of visible light for automatic landmine detection. *Advanced Sciences and Technologies for Security Applications*, 3:73–90, 2007.
- [58] JGM Schavemaker, F. Cremer, K. Schutte, and E. Den Breejen. Infrared processing and sensor fusion for anti-personnel land-mine detection. *Proceedings of IEEE Student Branch Eindhoven: Symposium Imaging*, page 61–71, 2000.
- [59] Frank Cremer and Wim de Jong. Infrared polarization measurements and modeling applied to surface-laid antipersonnel landmines. *Optical Engineering*, 41:1021, 2002.
- [60] Wilhelmus A. C. M. Messelink, Klammer Schutte, Albert M. Vossepoel, Frank Cremer, John G. M. Schavemaker, and Eric den Breejen. Feature-based detection of landmines in infrared images. *Detection and Remediation Technologies for Mines and Minelike Targets VII*, 4742:108–119, 2002.
- [61] Peter Ngan, Sigberto A. Garcia, Eugene L. Cloud, Herbert A. Duvoisin III, Daniel T. Long, and Jay K. Hackett. Development of automatic target recognition for infrared sensor-based close-range land mine detector. *Detection Technologies for Mines and Minelike Targets*, 2496:881–889, 6 1995.
- [62] Maria Garcia-Fernandez, Yuri Alvarez-Lopez, and Fernando Las Heras. Autonomous airborne 3d sar imaging system for subsurface sensing: Uwb-gpr on board a uav for landmine and ied detection. *Remote Sensing*, 11(20):2357, 2019.
- [63] Timothy De Smet, Alex Nikulin, William Frazer, Jasper Baur, Jacob Abramowitz, Daniel Finan, Sean Denara, Nicholas Aglietti, and Gabriel Campos. Drones and ‘butterflies’: A low-cost uav system for rapid detection and identification of unconventional minefields. *The Journal of Conventional Weapons Destruction*, 22(3):10, 2018.
- [64] Alex Nikulin, Timothy S De Smet, Jasper Baur, William D Frazer, and Jacob C Abramowitz. Detection and identification of remnant pfm-1 ‘butterfly mines’ with a uav-based thermal-imaging protocol. *Remote Sensing*, 10(11):1672, 2018.
- [65] Jasper Baur, Gabriel Steinberg, Alex Nikulin, Kenneth Chiu, and Timothy S de Smet. Applying deep learning to automate uav-based detection of scatterable landmines. *Remote Sensing*, 12(5):859, 2020.

- [66] Serkan Kaya and Ugur Murat Leloglu. Buried and surface mine detection from thermal image time series. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 10:4544–4552, 2017.
- [67] Yuvaraj Ganesh, Ramya Raju, and Rajeshwari Hegde. Surveillance drone for landmine detection. *Proceedings - 2015 21st Annual International Conference on Advanced Computing and Communications, ADCOM 2015*, pages 33–38, 8 2016.
- [68] Jordan Milton Malof, Loren Nolte, and Rebecca Willett. Exploiting multi-look information for landmine detection in forward looking infrared. Duke University, 2013. PhD Dissertation.
- [69] Ljuan L. Gurdev, Dimitar V. Stoyanov, Tanja N. Dreischuh, Christo N. Protochristov, and Orlin I. Vankov. Gamma-ray backscattering tomography approach based on the lidar principle. *IEEE Transactions on Nuclear Science*, 54:262–275, 2 2007.
- [70] Tanja N Dreischuh, Ljuan L Gurdev, Dimitar V Stoyanov, Christo N Protochristov, and Orlin I Vankov. Application of a lidar-type gamma-ray tomography approach for detection and identification of buried plastic landmines. In *14th International School on Quantum Electronics: Laser Physics and Applications*, volume 6604, pages 485–489. SPIE, 2007.
- [71] H Meurer²⁸, M Wehner²⁹, S Schillberg³⁰, K Hund-Rinke³¹, Ch Kühn³¹, and N Raven³¹. An emerging remote sensing technology and its potential impact on mine action. In *International Symposium” Humanitarian Demining 2010*, page 66, 2010.
- [72] Velodyne. Velodyne’s guide to lidar wavelengths, 2018. <https://velodynelidar.com/blog/guide-to-lidar-wavelengths/>.
- [73] IRTC. Technology overview frequently asked questions about wide-area assessment for munitions response projects, 2010.
- [74] Vallon GmbH. Vmr3g – dual-sensor detector, 2020. <https://www.vallon.de/en/dual-sensor-detectors-gpr/vmr3g-dual-sensor-detector-minehound>.
- [75] GapEOD. Ultratem, 2020.
- [76] Geometrics. Metalmapper 2x2 advanced time domain em uxo sensor. 2018.
- [77] E. Gasperikova, J. Torquil Smith, H. Frank Morrison, and A. Becker. Berkeley uxo discriminator (bud). European Association of Geoscientists Engineers, 2007.
- [78] Karl N. Kappler and Erika Gasperikova. A hybrid method for uxo vs. non-uxo discrimination. *Journal of Environmental and Engineering Geophysics*, 16:177–189, 2011.
- [79] James B Kingdon, Bruce J Barrow, Thomas H Bell, David C George, Glenn R Harbaugh, and Daniel A Steinhurst. Temtads adjunct sensor systems-hand-held emi sensor for cued uxo discrimination (estcp mr-200807) and man-portable emi array for uxo detection and discrimination (estcp mr-200909) final report. *US Naval Research Laboratory Report NRL/MR/6110-12-9401*, 2012.

- [80] Juan Pablo Fernández, Benjamin E Barrowes, Tomasz M Grzegorzczak, Nicolas Lhomme, Kevin O'Neill, and Fridon Shubitidze. A man-portable vector sensor for identification of unexploded ordnance. *IEEE Sensors Journal*, 11(10):2542–2555, 2011.
- [81] Zhicheng Yang, Shenggang Yan, Lianping Chen, and Bin Li. Ferromagnetic object localization based on improved triangulating and ranging. *IEEE Magnetics Letters*, 10:1–5, 2019.
- [82] Dingchen Wang, Songrui Wei, Anran Yuan, Fanghua Tian, Kaiyan Cao, Qizhong Zhao, Yin Zhang, Chao Zhou, Xiaoping Song, Dezhen Xue, et al. Machine learning magnetic parameters from spin configurations. *Advanced Science*, 7(16):2000566, 2020.
- [83] Bülent Oruç. Location and depth estimation of point-dipole and line of dipoles using analytic signals of the magnetic gradient tensor and magnitude of vector components. *Journal of Applied Geophysics*, 70(1):27–37, 2010.
- [84] Jan Francke and Alexey Dobrovolskiy. Challenges and opportunities with drone-mounted gpr. In *First International Meeting for Applied Geoscience & Energy*, pages 3043–3047. Society of Exploration Geophysicists, 2021.
- [85] Lei Xu, Ning Zhang, Liqing Fang, Huadong Chen, Pengfei Lin, and Chunsheng Lin. Simulation analysis of magnetic gradient full-tensor measurement system. *Mathematical Problems in Engineering*, pages 1–13, 2021.
- [86] Heidi Myers, Daniel Lathrop, and Vedran Lekic. Vehicle-based anomaly detection using artificial intelligence and combined environmental and geophysical sensor data, Patent Pending. USPO - Application - UMD: PS-2021-076.
- [87] Hilal Ai Azry, Hussain Ai Zubaidy, S.S. Webster, W. Witham, and D. Isles. Exploration for cyprus style copper deposits, sultanate of oman: A case history. *Exploration Geophysics*, 24:315–321, 9 1993.
- [88] P J Gunn and M C Dentith. Magnetic responses associated with mineral deposits. *AGSO Journal of Australian Geology and Geophysics*, 17:145–158, 1997.
- [89] William J Hinze, Ralph R B Von Frese, R Von Frese, and Afif H Saad. *Gravity and magnetic exploration: Principles, practices, and applications*. Cambridge University Press, 2013.
- [90] Hashem Shahsavani. Comparison of a low-cost magneto-inductive magnetometer with a proton magnetometer: a case study on the galali iron ore deposit in western iran. *Near Surface Geophysics*, 17:69–84, 2019.
- [91] F Accomando, A Bonfante, M Buonanno, G Florio, J Natale, and S Vitale. Mapping an igneous dike in carbonate rocks by drone-borne magnetometry. volume 2022, pages 1–5. EAGE Publications BV, 2022.

- [92] Marina Díaz Michelena, Rolf Kilian, Miguel Ángel Rivero, Sergio Fernández Romero, Francisco Ríos, José Luis Mesa, and Andrés Oyarzún. Magnetometric surveys for the non-invasive surface and subsurface interpretation of volcanic structures in planetary exploration, a case study of several volcanoes in the iberian peninsula. *Remote Sensing* 2022, Vol. 14, Page 2039, 14:2039, 4 2022.
- [93] Yoshifumi Nogi, John A Tarduno, and William W Sager. Inferences about the nature and origin of basalt sequences from the cretaceous mid-pacific mountains (sites 865 and 866), as deduced from downhole magnetometer logs. *Proceedings of the Ocean Drilling Program, Scientific Results*, 143:381–388, 1995.
- [94] Anderson Costa Santos, Webster Ueipass Mohriak, Mauro Cesar Geraldles, Werlem Holanda Santos, Cosme Ferreira Ponte-Neto, and Natasha Stanton. Compiled potential field data and seismic surveys across the eastern brazilian continental margin integrated with new magnetometric profiles and stratigraphic configuration for trindade island, south atlantic, brazil. *International Geology Review*, 61:1728–1744, 2019.
- [95] C Barrère, J Ebbing, and L Gernigon. 3-d density and magnetic crustal characterization of the southwestern barents shelf: implications for the offshore prolongation of the norwegian caledonides. *Geophysical Journal International*, 184:1147–1166, 2011.
- [96] F. S. Birch. Bedrock depth estimates from ground magnetometer profiles. *Groundwater*, 22:427–432, 7 1984.
- [97] Lindrith Cordell and V J S Grauch. *Mapping basement magnetization zones from aeromagnetic data in the San Juan Basin, New Mexico*. Society of Exploration Geophysicists, 1985.
- [98] Arne Døssing, H Ruth Jackson, Jürgen Matzka, Indrii Einarsson, Thorkild Maack Rasmussen, Arne Vestergaard Olesen, and JM Brozena. On the origin of the amerasia basin and the high arctic large igneous province—results of new aeromagnetic data. *Earth and Planetary Science Letters*, 363:219–230, 2013.
- [99] Filippo Accomando, Andrea Vitale, Antonello Bonfante, Maurizio Buonanno, and Giovanni Florio. Performance of two different flight configurations for drone-borne magnetic data. *Sensors*, 21:5736, 8 2021.
- [100] Yunes Sh Alqudsi, Abdulhameed S Alsharafi, and Abdulwahid Mohamed. A review of airborne landmine detection technologies: Unmanned aerial vehicle-based approach. *2021 International Congress of Advanced Technology and Engineering (ICOTEN)*, pages 1–5, 2021.
- [101] Lee Sun Yoo, Jung Han Lee, Sung Hyub Ko, Seom Kyu Jung, Seung Hun Lee, and Yong Kuk Lee. A drone fitted with a magnetometer detects landmines. *IEEE Geoscience and Remote Sensing Letters*, 17:2035–2039, 12 2020.
- [102] R Blake. Potential theory in gravity magnetic application, 1995.

- [103] Joanna Hamlyn. *Magnetometry method*. CRC Press, 2022.
- [104] Nicolas Florsch, Christian Camerlynck, Michel Kammenthaler, and Frédéric Muhlach. *Everyday Applied Geophysics 2: Electromagnetics and Magnetism*. ISTE Press Limited-Elsevier Incorporated, 2018.
- [105] Ph D Nikulin, Ph D deSmet, Andrii Puliaiev, Vasyl Zhurakhov, Sofia Fasullo, Gabriel Chen, Isaac Spiegel, and Kaylee Cappuccio. Automated uas aeromagnetic surveys to detect mbrl unexploded ordnance. *The Journal of Conventional Weapons Destruction*, 24:13, 2020.
- [106] Manuel Ricardo Pérez Cerquera, Julian David Colorado Montaña, and Iván Mondragón. Uav for landmine detection using sdr-based gpr technology. *Robots Operating in Hazardous Environments*, 12 2017.
- [107] Bruno Gavazzi, Pauline Le Maire, Marc Munsch, and Aline Dechamp. Fluxgate vector magnetometers: A multisensor device for ground, uav, and airborne magnetic surveys. *Leading Edge*, 35:795–797, 9 2016.
- [108] Samuel Prager and Mahta Moghaddam. Application of ultra-wideband synthesis in software defined radar for uav-based landmine detection. *International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 10115–10118, 2019.
- [109] Mioara Manda and · Aude Chambodut. Geomagnetic field processes and their implications for space weather. *Surveys in Geophysics*, 41:1611–1627, 2020.
- [110] Karl N. Kappler, Daniel D. Schneider, Laura S. MacLean, and Thomas E. Bleier. Identification and classification of transient pulses observed in magnetometer array data by time-domain principal component analysis filtering. *Earthquake Science*, 30:193–207, 8 2017.
- [111] Clive Foss. Magnetic data enhancements and depth estimation. *Encyclopedia of solid earth geophysics*, pages 1–14, 2020.
- [112] R P Riddihough. Diurnal corrections to magnetic surveys—an assessment of errors. *Geophysical Prospecting*, 19:551–567, 1971.
- [113] Harold L Yarger, Robert R Robertson, and Robert L Wentland. Diurnal drift removal from aeromagnetic data using least squares. *Geophysics*, 43:1148–1156, 1978.
- [114] Marc A Vallée, James Craven, Larry Newitt, Pierre Keating, Régis Dumont, and Ian Ferguson. Guidelines for location and use of base station data in aeromagnetic survey processing. OnePetro, 2006.
- [115] Leo James Peters. The direct approach to magnetic interpretation and its practical application. *Geophysics*, 14:290–320, 1949.
- [116] R A Smith. Some depth formulae for local magnetic and gravity anomalies. *Geophysical Prospecting*, 7:55–63, 1959.

- [117] M H P Bott and R A Smith. The estimation of the limiting depth of gravitating bodies. *Geophysical Prospecting*, 6:1–10, 1958.
- [118] Richard J Blakely. *Potential theory in gravity and magnetic applications*. Cambridge university press, 1996.
- [119] Ab B Reid, J M Allsop, Harald Granser, A J t Millett, and I W Somerton. Magnetic interpretation in three dimensions using euler deconvolution. *Geophysics*, 55:80–91, 1990.
- [120] D T Thompson. Euldp: A new technique for making computer-assisted depth estimates from magnetic data. *Geophysics*, 47:31–37, 1982.
- [121] L A Gallardo, S L Fontes, M A Meju, M P Buonora, and P P De Lugao. Robust geophysical integration through structure-coupled joint inversion and multispectral fusion of seismic reflection, magnetotelluric, magnetic, and gravity images: Example from santos basin, offshore brazil. *Geophysics*, 77:B237–B251, 2012.
- [122] Emad Ghalenoei, Jan Dettmer, Mohammed Y Ali, and Jeong Woo Kim. Gravity and magnetic joint inversion for basement and salt structures with the reversible-jump algorithm. *Geophysical Journal International*, 227:746–758, 2021.
- [123] Hermann Zeyen and Jaume Pous. 3-d joint inversion of magnetic and gravimetric data with a priori information. *Geophysical Journal International*, 112:244–256, 1993.
- [124] Stephen D Billings, Leonard R Pasion, and Douglas W Oldenburg. Discrimination and identification of uxo by geophysical inversion. phase ii: Inversion of total-field magnetics, 2002.
- [125] Leonard R Pasion, Stephen D Billings, and Douglas W Oldenburg. Joint and cooperative inversion of magnetic and time domain electromagnetic data for the characterization of uxo. pages cp–190. European Association of Geoscientists Engineers, 2003.
- [126] Huan Liu, Haobin Dong, Jian Ge, and Zheng Liu. An overview of sensing platform-technological aspects for vector magnetic measurement: A case study of the application in different scenarios. *Measurement*, 187:110352, 2022.
- [127] Hongfeng Pang, Shitu Luo, Qi Zhang, Ji Li, Dixiang Chen, Mengchun Pan, and Feilu Luo. Calibration of a fluxgate magnetometer array and its application in magnetic object localization. *Measurement Science and Technology*, 24, 2013.
- [128] J. M.G. Merayo, P. Brauer, F. Primdahl, J. R. Petersen, and O. V. Nielsen. Scalar calibration of vector magnetometers. *Measurement Science and Technology*, 11:120–132, 2000.
- [129] Nils Olsen, Fritz Primdahl, Lars Tøffner-Clausen, Terence Sabaka, Peter Brauer, Jose Merayo, John Jørgensen, Otto Nielsen, Fritz Primdahl, J. M. Léger, and Torben Risbo. Calibration of the Ørsted vector magnetometer. *Earth, Planets and Space*, 55:11–18, 2003.
- [130] Timo Pylvänäinen. Automatic and adaptive calibration of 3d field sensors. *Applied Mathematical Modelling*, 32(4):575–587, 2008.

- [131] Phil W. Schmidt and David A. Clark. The magnetic gradient tensor: Its properties and uses in source characterization. *Leading Edge (Tulsa, OK)*, 25:75–78, 4 2006.
- [132] Long Pham and Anthony Desimone. Vibration rectification in mems accelerometers. *Analog Devices, Inc.*, 2017.
- [133] P.W. Schmidt and D.A. Clark. Advantages of measuring the magnetic gradient tensor. *Preview*, 85:26–30, 2000.
- [134] Hongpeng Wang, Tao Meng, Wang Luo, Ruiping Yang, Huan Liu, Jian Ge, and Haobin Dong. Overhauser sensor array based 3-d magnetic gradiometer for the detection of shallow subsurface unexploded ordnance. *IEEE Transactions on Instrumentation and Measurement*, 72:1–11, 2023.
- [135] Huanghuang Jin, Jian Guo, Hongbo Wang, Zhihong Zhuang, Jie Qin, and Tonglei Wang. Magnetic anomaly detection and localization using orthogonal basis of magnetic tensor contraction. *IEEE Transactions on Geoscience and Remote Sensing*, 58(8):5944–5954, 2020.
- [136] R. Wiegert and J. Oeschger. Generalized magnetic gradient contraction based method for detection, localization and discrimination of underwater mines and unexploded ordnance. In *Proceedings of OCEANS 2005 MTS/IEEE*, pages 1325–1332 Vol. 2, 2005.
- [137] Gene H Golub and Charles F Van Loan. *Matrix computations*. JHU press, 2013.
- [138] L. B. Pedersen and T. M. Rasmussen. The gradient tensor of potential field anomalies: Some implications on data collection and data processing of maps. *GEOPHYSICS*, 55:1558–1566, 12 1990.
- [139] David A Clark. Corrigendum to: New methods for interpretation of magnetic vector and gradient tensor data i: eigenvector analysis and the normalised source strength. *Exploration Geophysics*, 45:324, 2014.
- [140] Majid Beiki, David A Clark, James R Austin, and Clive A Foss. Estimating source location using normalized magnetic source strength calculated from magnetic gradient tensor datanormalized source strength. *Geophysics*, 77:J23–J37, 2012.
- [141] David A Clark. New methods for interpretation of magnetic vector and gradient tensor data ii: application to the mount leyshon anomaly, queensland, australia. *Exploration Geophysics*, 44:114–127, 2013.
- [142] W. E. Doll, J. E. Nyquist, L. P. Beard, and T. J. Gamey. Airborne geophysical surveying for hazardous waste site characterization on the oak ridge reservation, tennessee. *GEO-PHYSICS*, 65:1372–1387, 9 2000.
- [143] T Jeffrey Gamey, J Scott Holladay, and Reinhold Mahler. Airborne measured analytic signal for uxo detection. pages cp–204. European Association of Geoscientists Engineers, 1997.

- [144] D W Smellie. Elementary approximations in aeromagnetic interpretation. *Geophysics*, 21:1021–1040, 1956.
- [145] F. Rosenblatt. The perceptron, a perceiving and recognition automation, 1957.
- [146] D. Rumelhart, G. Hinton, and J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.
- [147] Grégoire Montavon, Geneviève Orr, and Klaus-Robert Müller. *Neural networks: tricks of the trade*, volume 7700. springer, 2012.
- [148] A. Gupta. A comprehensive guide to optimizers in deep learning. <https://www.analyticsvidhya.com/blog/2021/10/a-comprehensive-guide-on-deep-learning-optimizers/>.
- [149] S. Kotsiantis, B. Zaharakis, and P. Pintelas. Machine learning: a review of classification and combining techniques. *Artificial Intelligence Review*, 26:159–190, 2006.
- [150] U. Acharya, S. Oh, J. Hagiwara, M. Tan, A. Adam, A. Gertych, and R. Tan. A deep convolutional neural network model to classify heartbeats. *Computers in Biology and Medicine*, 89:389–396, 2017.
- [151] H. Ramchoun, M. Idrissi, Y. Ghanou, and M. Ettaiouil. Multilayer perceptron: Architecture optimization and training. *International Journal of Interactive Multimedia and Artificial Intelligence*, 4:26–30, 2016.
- [152] S. Albawi, T. Mohammed, and S. Al-Zawi. Understanding of a convolutional neural network. 2017.
- [153] Juan Pablo Fernández, Keli Sun, Benjamin Barrowes, Kevin O’Neill, Irma Shamatava, Fridon Shubitidze, and Keith D Paulsen. Inferring the location of buried uxo using a support vector machine. In *Detection and Remediation Technologies for Mines and Minelike Targets XII*, volume 6553, pages 59–67. SPIE, 2007.
- [154] Minh-Tan Pham and Sébastien Lefèvre. Buried object detection from b-scan ground penetrating radar data using faster-rcnn. In *IGARSS 2018-2018 IEEE international geoscience and remote sensing symposium*, pages 6804–6807. IEEE, 2018.
- [155] Xavier Núñez-Nieto, Mercedes Solla, Paula Gómez-Pérez, and Henrique Lorenzo. Gpr signal characterization for automated landmine and uxo detection based on machine learning techniques. *Remote sensing*, 6(10):9729–9748, 2014.
- [156] Francois Chollet. *Deep learning with Python*. Simon and Schuster, 2021.
- [157] Pierre Raybaut. Spyder-documentation. Available online at: pythonhosted.org, 2009.
- [158] Dwain K Butler, Janet E Simms, John S Furey, and Hollis H Bennett. Review of magnetic modeling for uxo and applications to small items and close distances. *Journal of Environmental and Engineering Geophysics*, 17(2):53–73, 2012.

- [159] Mick Emil Kolster, Mark David Wigh, Eduardo Lima Simões da Silva, Tobias Bjerg Vilhelmsen, and Arne Døssing. High-speed magnetic surveying for unexploded ordnance using uav systems. *Remote Sensing*, 14(5):1134, 2022.
- [160] Jiazeng Wang, Ying Shen, Rui Zhao, Chunkai Zhou, and Junqi Gao. Estimation of dipole magnetic moment orientation based on magnetic signature waveform analysis by a magnetic sensor. *Journal of Magnetism and Magnetic Materials*, 505:166761, 2020.
- [161] Mark David Wigh, Mick Emil Kolster, Thomas Mejer Hansen, and Arne Døssing. Classification of uxo and non-uxo from magnetic anomaly data: a case study on inversion of drone magnetic data from rømø, denmark. *Geophysical Journal International*, 234(2):915–932, 2023.
- [162] Roy Wiegert and John Oeschger. Generalized magnetic gradient contraction based method for detection, localization and discrimination of underwater mines and unexploded ordnance. In *Proceedings of OCEANS 2005 MTS/IEEE*, pages 1325–1332. IEEE, 2005.
- [163] Takaaki Nara, Satoshi Suzuki, and Shigeru Ando. A closed-form formula for magnetic dipole localization by measurement of its magnetic field and spatial gradients. *IEEE transactions on magnetics*, 42(10):3291–3293, 2006.
- [164] Yijie Qin, Maobin Li, Keyan Li, Yang Pan, Xiaofei Yang, and Jun Ouyang. Target magnetic moment orientation estimation method based on full magnetic gradient orthonormal basis function. *AIP Advances*, 12(3), 2022.
- [165] Valeria Paoletti, Antonio Buggi, and Roman Pašteka. Uxo detection by multiscale potential field methods. *Pure and Applied Geophysics*, 176:4363–4381, 2019.
- [166] Volkmar Schmidt, Michael Becken, and Jörg Schmalzl. A uav-borne magnetic survey for archaeological prospection of a celtic burial site. *First Break*, 38(8):61–66, 2020.
- [167] J Bradley Nelson. Calculation of the magnetic gradient tensor from total field gradient measurements and its application to geophysical interpretation. *Geophysics*, 53:957–966, 1988.
- [168] Yongjun Wang, Zhi Li, and Xiang Li. External disturbances rejection for vector field sensors in attitude and heading reference systems. *Micromachines*, 11(9):803, 2020.
- [169] Jialong Dai, Lifei Meng, Kui Huang, Chaoqun Xu, Zhong Yi, Qi Xiao, Chaobo Liu, Bin Wang, Chao Zhang, Shaohua Zhang, et al. Research on magnetic characteristics of small uav for aeromagnetic measurement. In *IOP Conference Series: Earth and Environmental Science*, volume 310, page 032068. IOP Publishing, 2019.
- [170] Julien Metge, Rémi Megret, Audrey Giremus, Yannick Berthoumieu, and Christophe Mazel. Dynamic magnetic field compensation for micro uav attitude estimation. In *2013 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 725–733. IEEE, 2013.
- [171] David Clarke. Magnetic signature of brushless electric motors. 2006.

- [172] Yaxin Mu, Xiaojuan Zhang, Wupeng Xie, and Yaoxin Zheng. Automatic detection of near-surface targets for unmanned aerial vehicle (uav) magnetic survey. *Remote Sensing*, 12, 2020.
- [173] N.R. Syed. How to read a dc motor datasheet, 2018. <https://nrsyed.com/2018/01/21/how-to-read-a-dc-motor-datasheet/>.
- [174] Jonathan Law and Richard Rennie. *A dictionary of physics*. OUP Oxford, 2015.
- [175] Salih Baris Ozturk and Hamid A Toliyat. Direct torque and indirect flux control of brushless dc motor. *IEEE/ASME Transactions On Mechatronics*, 16(2):351–360, 2010.
- [176] Luca Zaccarian. Dc motors: dynamic model and control techniques. *Lecture notes., roma, italy*, 2012.
- [177] Rithvik Gambhir and Akshay Kumar Jha. Brushless dc motor: Construction and applications. *Int. J. Eng. Sci*, 2(5):72–77, 2013.
- [178] AeroVironment, 2024. <https://www.avinc.com/ugv/teodor-evo>.
- [179] Julián Colorado, M Perez, I Mondragon, Diego Mendez, Carlos Parra, C Devia, J Martinez-Moritz, and L Neira. An integrated aerial system for landmine detection: Sdr-based ground penetrating radar onboard an autonomous drone. *Advanced Robotics*, 31(15):791–808, 2017.
- [180] Manuel Ricardo Pérez Cerquera, Julian David Colorado Montaña, Iván Mondragón, and H Canbolat. Uav for landmine detection using sdr-based gpr technology. *Robots Operating in Hazardous Environments*, pages 26–55, 2017.
- [181] Carlo Noviello, Gianluca Gennarelli, Giuseppe Esposito, Giovanni Ludeno, Giancarmine Fasano, Luigi Capozzoli, Francesco Soldovieri, and Ilaria Catapano. An overview on down-looking uav-based gpr systems. *Remote Sensing*, 14(14):3245, 2022.
- [182] Florin-Lucian Chiper, Alexandru Martian, Calin Vladeanu, Ion Marghescu, Razvan Craciunescu, and Octavian Fratu. Drone detection and defense systems: Survey and a software-defined radio-based solution. *Sensors*, 22(4):1453, 2022.
- [183] Daniel P Maxbauer, Joshua M Feinberg, and David L Fox. Magnetic mineral assemblages in soils and paleosols as the basis for paleoprecipitation proxies: A review of magnetic methods and challenges. *Earth-Science Reviews*, 155:28–48, 2016.
- [184] Tilo Von Dobeneck, Maximilian Müller, Benjamin Bosbach, and Andreas Klügel. Ground magnetic surveying and susceptibility mapping across weathered basalt dikes reveal soil creep and pedoturbation. *Frontiers in Earth Science*, 8:592986, 2021.
- [185] Susan L Brantley, Marina I Lebedeva, Victor N Balashov, Kamini Singha, Pamela L Sullivan, and Gary Stinchcomb. Toward a conceptual model relating chemical reaction fronts to water flow paths in hills. *Geomorphology*, 277:100–117, 2017.

- [186] Karianne J Bergen, Paul A Johnson, Maarten V de Hoop, and Gregory C Beroza. Machine learning for data-driven discovery in solid earth geoscience. *Science*, 363(6433):eaau0323, 2019.
- [187] Işık Yilmaz and Oguz Kaynar. Multiple regression, ann (rbf, mlp) and anfis models for prediction of swell potential of clayey soils. *Expert systems with applications*, 38(5):5958–5966, 2011.
- [188] Eugène C Ezin, Flora Giudicepietro, Simona Petrosino, Silvia Scarpetta, and Amalia Vana-core. Automatic discrimination of earthquakes and false events in seismological recording for volcanic monitoring. In *Neural Nets: 13th Italian Workshop on Neural Nets, WIRN VI-ETRI 2002 Vietri sul Mare, Italy, May 30–June 1, 2002 Revised Papers 13*, pages 140–145. Springer, 2002.
- [189] Yousef Asgari Nezhad, Ali Moradzadeh, and Mohammad Reza Kamali. A new approach to evaluate organic geochemistry parameters by geostatistical methods: A case study from western australia. *Journal of Petroleum Science and Engineering*, 169:813–824, 2018.
- [190] Haibin Di, Muhammad Shafiq, and Ghassan AlRegib. Patch-level mlp classification for improved fault detection. In *SEG Technical Program Expanded Abstracts 2018*, pages 2211–2215. Society of Exploration Geophysicists, 2018.
- [191] Karl Pearson. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11):559–572, 1901.
- [192] Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.
- [193] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [194] Jinkun Cheng, Ke Chen, and Mauricio D Sacchi. Robust principal component analysis (rpca) for seismic data denoising. *min J*, 1:2 μ , 2015.
- [195] MF Gazley, KS Collins, J Roberston, BR Hines, LA Fisher, and A McFarlane. Application of principal component analysis and cluster analysis to mineral exploration and mine geology. In *AusIMM New Zealand branch annual conference*, volume 2015, pages 131–139. Dunedin New Zealand, 2015.
- [196] Arthur Nanni, Ari Roisenberg, Maria Paula Casagrande Marimon, and Antonio Pedro Viero. The hydrochemical facies and anomalous fluoride content in the serra geral aquifer system, southern brazil: a gis approach with tectonic and principal component analysis. *Environmental Geology*, 58:1247–1255, 2009.
- [197] Leonardo Tomazeli Duarte, Everton Z Nadalin, Kenji Nose Filho, Ricardo Antonio Zanetti, João MT Romano, and Martin Tygel. Seismic wave separation by means of robust principal component analysis. In *2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*, pages 1494–1498. IEEE, 2012.

- [198] Shuwen Yang, Yikun Li, Guangsheng Feng, and Lifeng Zhang. A method aimed at automatic landslide extraction based on background values of satellite imagery. *International Journal of Remote Sensing*, 35(6):2247–2266, 2014.
- [199] Zachary E Ross, Men-Andrin Meier, and Egill Hauksson. P wave arrival picking and first-motion polarity determination with deep learning. *Journal of Geophysical Research: Solid Earth*, 123(6):5120–5129, 2018.
- [200] Daigo Shoji, Rina Noguchi, Shizuka Otsuki, and Hideitsu Hino. Classification of volcanic ash particles using a convolutional neural network and probability. *Scientific reports*, 8(1):8111, 2018.
- [201] Mahdi Panahi, Nitheshnirmal Sadhasivam, Hamid Reza Pourghasemi, Fatemeh Rezaie, and Saro Lee. Spatial prediction of groundwater potential mapping based on convolutional neural network (cnn) and support vector regression (svr). *Journal of Hydrology*, 588:125033, 2020.
- [202] Syamil Mohd Razak and Behnam Jafarpour. Convolutional neural networks (cnn) for feature-based model calibration under uncertain geologic scenarios. *Computational Geosciences*, 24(4):1625–1649, 2020.
- [203] Xuejia Sang, Linfu Xue, Xiangjin Ran, Xiaoshun Li, Jiwen Liu, and Zeyu Liu. Intelligent high-resolution geological mapping based on slic-cnn. *ISPRS International Journal of Geo-Information*, 9(2):99, 2020.
- [204] Joakim Andén and Stéphane Mallat. Scattering representation of modulated sounds. *15th DAFX*, 9:17–21, 2012.
- [205] S. Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65:1331–1938, 2012.
- [206] Joan Bruna and Stephane Mallat. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:1872–1886, 2013.
- [207] Oren Rippel, Jasper Snoek, and Ryan P Adams. Spectral representations for convolutional neural networks. *arXiv Preprint*, 2015.
- [208] Seda Arslan Tuncer, Hakan Ayyıldız, Mehmet Kalaycı, and Taner Tuncer. Scat-net: Covid-19 diagnosis with a cnn model using scattergram images. *Computers in Biology and Medicine*, 135:104579, 2021.
- [209] Yash Bhalgat, Jean Charléty, and Laurent Duval. Catseyes: categorizing seismic structures with tessellated scattering wavelet networks. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1618–1622. IEEE, 2018.
- [210] Ben Bougher and Felix J Herrmann. Using the scattering transform to predict stratigraphic units from well logs. *CSEG Recorder*, 41(1):22–25, 2016.

- [211] Leonard Seydoux, Piero Poli, Maarten de Hoop, and Michel Campillo. Detection and clustering of seismic events in continuous seismic data with scattering transform. In *Geophysical Research Abstracts*, volume 21, 2019.
- [212] T. Dayanand, 2020. Geological Image Similarity, kaggle.com.
- [213] Keras.io, 2015. Keras Software.
- [214] Laurens Van Der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.