

Field Calibration of Hyperspectral Cameras for Terrain Inference

Nathaniel Hanson^{1,2*}, Benjamin Pyatski¹, Samuel Hibbard¹, Gary Lvov¹, Oscar De La Garza¹, Charles DiMarzio¹, Kristen L. Dorsey¹, and Taşkın Padır^{1,3}

Abstract—Intra-class terrain differences such as water content directly influence a vehicle’s ability to traverse terrain, yet RGB vision systems may fail to distinguish these properties. Evaluating a terrain’s spectral content beyond red-green-blue wavelengths to the near infrared spectrum provides useful information for intra-class identification. However, accurate analysis of this spectral information is highly dependent on ambient illumination. We demonstrate a system architecture to collect and register multi-wavelength, hyperspectral images from a mobile robot and describe an approach to reflectance calibrate cameras under varying illumination conditions. To showcase the practical applications of our system, HYPER DRIVE, we demonstrate the ability to calculate vegetative health indices and soil moisture content from a mobile robot platform.

Index Terms—Field Robots; Robotics and Automation in Agriculture and Forestry; Calibration and Identification

I. INTRODUCTION

UNDERSTANDING *where* to drive is a critical question for autonomous vehicles [1]. Many systems operate by identifying and defining safe regions for traversability such as road, grass, soil, and sand from RGB camera images [2]. While broad terrain labels are useful in semantic segmentation, they do not capture intra-class differences—an oil slick on a road, ice on soil, or water-logged grass—that may impact vehicle traversability. Adding spectral data as a complement to RGB images improves class separation [3], so incorporating intra-class variations into robot intelligence may similarly improve performance. To realize this vision—dynamically identifying these differences onboard a field robot—new sensing and algorithmic approaches are required.

Hyperspectral imaging (HSI) is well-established in remote sensing for airborne terrain monitoring, material detection [4],

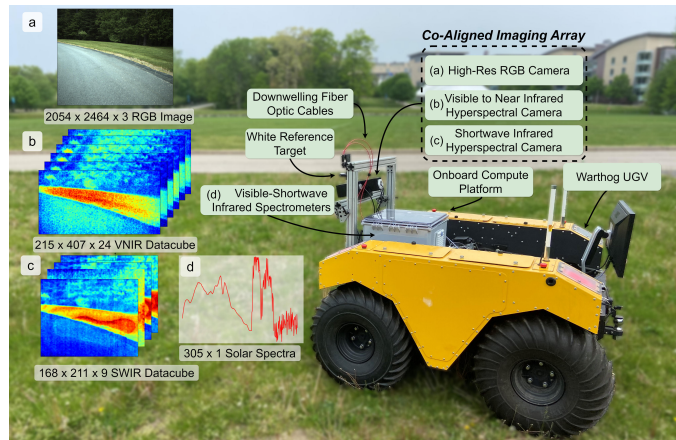


Fig. 1: Our hyperspectral imaging system, HYPER DRIVE, mounted to an off-road mobile robot. Images are collected from (a) a high resolution RGB camera, (b) a visible-near infrared (VNIR) hyperspectral camera, and (c) a shortwave infrared (SWIR) hyperspectral camera. We use combined (d) point spectrometers to measure a white-reference signal and calibrate the hyperspectral cameras.

and anomaly identification [5]. In contrast to the three intensity values present in RGB cameras, hyperspectral imagers record dozens of discrete wavelengths. Characteristic absorbance peaks across the visible-to-infrared spectrum due to water, organic compounds, and other chemical compositions will then provide classification insight. While HSI often improves inter-class separability over RGB alone [3], its use has been constrained to environments with well-controlled illumination or frequent camera re-calibration using a target. These constraints are untenable for robots in the field.

We propose to use field robot-based HSI to identify intra-class terrain variations and dynamically calibrate changes in illumination conditions that would otherwise distort the spectral information. We apply HSI to identify intra-class terrain variations and demonstrate a system, HYPER DRIVE, to capture datacubes (i.e., a multi-dimensional view across spectrum and space) from a moving platform (Fig. 1). We develop and validate a data-driven approach to dynamically calibrate the imaging areas via a learned up-sampling network that references a calibration target continuously sampled by an on-board point spectrometer. Our work shows that forward-facing HSI enables low computational cost differentiation of visually similar terrain and is practical for field robotics.

This work contributes the following to the state-of-the-art:

- The design and development of HYPER DRIVE, a field robot specific hyperspectral camera system,
- A method for spectrometer-hyperspectral joint calibration under varying illumination conditions, and
- The evaluation of HYPERDRIVE’s performance in estimating soil moisture content and vegetative health.

Manuscript received: May 23, 2025; Revised August 12, 2025; Accepted September 16, 2025.

This paper was recommended for publication by Editor Lucia Pallottino upon evaluation of the Associate Editor and Reviewers’ comments.

This research was sponsored by the United States Army Corps of Engineers (USACE) Engineer Research and Development Center (ERDC) Geospatial Research Laboratory (GRL) and was accomplished under Cooperative Agreement Federal Award Identification Number (FAIN) W9132V-22-2-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of USACE EDRC GRL or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

¹Nathaniel Hanson, Benjamin Pyatski, Samuel Hibbard, Gary Lvov, Oscar De La Garza, Charles DiMarzio, Kristen L. Dorsey, and Taşkın Padır are with the Electrical and Computer Engineering Department, Northeastern University, Boston, Massachusetts, USA.

²Nathaniel Hanson is with the Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, Massachusetts, USA.

³Taşkın Padır holds concurrent appointments as a Professor of Electrical and Computer Engineering at Northeastern University and as an Amazon Scholar. This paper describes work performed at Northeastern University and is not associated with Amazon.

*Correspondence: nathaniel.hanson@ll.mit.edu

Digital Object Identifier (DOI): see top of this page.




Fig. 2: The likelihood of incident photon eliciting a digital count and sampling widths of the spectral sensors in HYPER DRIVE.

II. BACKGROUND

A. Hyperspectral imaging

Hyperspectral cameras measure light intensity at multiple wavelengths in the visible (400 – 800 nm), the visible to near infrared (VNIR, 400 – 1100 nm), and the shortwave infrared (SWIR, 1100 – 2500 nm) spectra. The material composition of a scene may be identified from HSI through the spectral absorbances of various materials at specific wavelengths (i.e., their spectral “signatures”). For intra-class terrain identification, wavelengths related to water absorption (970 nm [6], 1200 nm, and 1400 nm [7]), and chlorophyll concentration (730 nm, e.g., from plants), are critical.

B. Related work

Under controlled settings, hyperspectral sensors are calibrated to measure spectral radiance. This approach is less practical in dynamic and autonomous systems because measuring outdoor radiance requires knowing atmospheric conditions and the position of the sun [8]. To address this challenge, recent work has scaled spectral readings between spectroradiometers mounted on ground-based and unmanned aerial vehicles (UAVs) [9] and reconstructed reflectance through a subspace representation and periodic imaging of a target [10]. This approach requires a robot to always be within navigable distance of a calibration target or to carry a large calibration target onboard. [11] demonstrated on-robot calibration for a VNIR snapshot camera, but significant changes in the incident illumination required additional collection and calibration. This approach also assumes that spectral distributions are constant over the VNIR range [12], which does not hold in the SWIR range.

In autonomous systems, a more practical approach is to calculate the spectral reflectance, R , as a proportion of detected light relative to maximum signal intensity. Following this approach, an HSI-specific version of min-max normalization adapted from [13] is

$$\text{reflectance} = \frac{\text{signal} - \text{signal}_{\text{dark}}}{\text{signal}_{\text{reference}} - \text{signal}_{\text{dark}}} \quad (1)$$

where the dark signal is the reading when the camera is covered, and the reference signal is the camera reading from a target. Typically, compressed polytetrafluoroethylene (PTFE) is chosen as the reference target due to its uniform reflection of visible - SWIR radiation.

Several previous works have demonstrated HSI in robotics, albeit with assumptions about the respective camera’s orientation or working distance. [14] demonstrated hyperspectral imaging on a UAV through an initial spectral calibration. This approach is less feasible for ground vehicles due to the large relative motion of the sun and the presence of other light sources,

reflections, and objects on the ground. [15] demonstrated terrain classification by measuring the ground spectral signature near the tire contact point on a wheeled vehicle, but this approach required (i) a short distance (5 cm) between the spectrometer and the ground and (ii) active illumination. Integrating HSI across a range of autonomous vehicles will require large field of vision cameras rather than point spectrometers, dynamic and computationally feasible calibration strategies, and the use of only passive illumination. To the best of our knowledge, no work has addressed the challenge of dynamic calibration of hyperspectral cameras and generalizes across camera models and wavelengths.

III. SYSTEM ARCHITECTURE

A. Imaging System

This work employs the HSI system first developed in [16], which featured broader wavelength sensitivity than any current robot system and explored the use of non-visible light to understand broad inter-class terrain classes. The sensor array includes hyperspectral cameras and a high-resolution RGB camera paired with a down-welling spectrometer to measure reflected solar light. The VNIR camera (XIMEA SNAPSHOT NIR, IMEC) captures 660 – 900 nm using Fabry–Pérot filters across 24 spectral bands, while the SWIR camera (SNAPSHOT SWIR 9 Band, IMEC) captures 1100 – 1700 nm with 9 spectral bands. Combined, the system provides 33 channels over 1100 nm of the spectrum. The cameras are housed in a weatherized 3D-printed (Onyx, Markforged) casing with broadband transmissive windows (Gorilla Glass, Corning) for minimal light perturbation. The VNIR and SWIR hyperspectral cameras have fields of view (FOV) of 25°. The complete sensor package, including housing, measures 250 × 82 × 73 mm, and has a mass of \approx 5.0 kg. Data acquisition is handled by an onboard computer (Intel NUC 11, Core i7 Processor, 64 GB RAM, 2 TB SSD, NVIDIA RTX 2060 GPU), making the system practical for integration onto other UGVs.

The hyperspectral cameras are spatially registered with a 5 MP RGB machine vision camera with an FOV of 30.9°. A combined visible-shortwave infrared (VIS-SWIR) datacube is created using projective transforms and keypoint mapping. The final datacube (1012 × 1666 × 33) is cropped to overlapping regions and interpolated to match the RGB resolution.

B. Point Spectrometers

Complementing the cameras are two spectrometers covering 500 – 1700 nm. The VNIR spectrometer (Pebble VIS-NIR, Ibsen Photonics) captures 500–1100 nm with a silicon detector and 256 spectral pixels, and the SWIR spectrometer (Pebble NIR, Ibsen Photonics) spans 950 – 1700 nm with an InGaAs detector and 128 spectral pixels. Their overlapping range




Fig. 3: To calibrate a datacube, the workflow takes in raw HSI hypercube data and spectral data from the down welling spectrometer (yellow), compares it to static calibration references (black), transforms it (blue), and passes it to the relevant algorithms (red).

compensates for decreased efficiency at higher wavelengths and ensures enhanced spectral sensitivity. Fiber optic cables (Thor Labs) couple light into the spectrometers from ferrules offset 4 cm above a 99% Spectralon white reference target (Labsphere). The quantum efficiencies of the devices are plotted in Fig. 2. These spectrometer readings are used to dynamically generate reflectance calibrations according to the method discussed in Section IV.

IV. SPECTRAL WHITE REFERENCE GENERATION

A. Data Collection

To relate point spectrometer readings to hyperspectral images, we designed a set of experiments to measure the correspondence between the devices when they are viewing a calibration target made of the same material. The system was placed outdoors and aligned to the 5° azimuth with the solar track. This placement also ensured that no external shadows fell on either the spectrometer target or on the white reference target. As the solar track varies only slightly from one day to the next, we assumed that the maximum solar overhead angle was constant across the measurements.

Across three days, the system was operated with the parameters in Table I under low to full sunlight from early morning to early evening and varying precipitation conditions. Devices sensing in the SWIR domain typically have longer integration times than VNIR devices due to decreased photon energy and proportionally lower atmospheric transmission.

A synchronized sample was acquired from each device every 5 minutes and saved with the time stamp at which it was acquired. In total, more than 2500 samples were collected to form a set of calibration spectra P .

B. White Reference Generation Formulation

The terms in Table II define hyperspectral and spectrometer data terms that are relevant to our analysis. For all variables, the subscript indicates the type of device (spectrometer S or HSI I) and the wavelength range (VNIR or SWIR). In the general case, we desire to approximate a function f (2), which maps an S to I as:

$$I = f(S) \quad (2)$$

TABLE I: HYPER DRIVE Training Device Parameters

Parameter	Value
VNIR HSI Frame rate	10 Hz
VNIR HSI Integration Time	0.5 ms
SWIR HSI Frame rate	10 Hz
SWIR HSI Integration Time	1.0 ms
VNIR Spectrometer Integration Time	0.5 ms
SWIR Spectrometer Integration Time	50 ms

TABLE II: Targetless White Reference Generation Parameters

Parameter	Description
t	Integration time, the time for which the sensor was exposed to light.
$I\{H, W, \Lambda_I\}$	3D hyperspectral datacube with height H , width W , and wavelength channels Λ in digital counts. Pixels of the datacube are indexed using the column c , row r , and channel λ . $I \in \mathbb{Z}^{H \times W \times \Lambda}$
$S\{\Lambda_S\}$	1-D Vector, representing the digital counts of the spectrometer. $S \in \mathbb{Z}^{\Lambda_S}$
D	Bit depth of a device, where the subscript represents the specific device and range
P	Set of all observed spectra
Λ	Set of wavelengths, subscript indicates the specific device or subset

As we have sets of devices that cover the wavelength spectrum, we consider the calibration of VNIR and SWIR devices separately. In this calibration routine, we assume integration times are set low enough to avoid saturation. At every time step, we assume there is a spectrometer reading of digital counts S_{raw} and a hyperspectral datacube also of digital counts I_{raw} . The saturation point of the digital count for these devices varies widely by device; therefore, we first begin by scaling each measurement to be normalized in the range $[0, 1]$ by dividing by the maximum digital count value at saturation. We denote by $S \in \mathbb{Z}^+$ and $I \in \mathbb{Z}^+$ the raw, unnormalized spectrometer and hyperspectral measurements in digital counts. The normalized quantities, S_{norm} and I_{norm} , are obtained by scaling the raw counts to the range $[0, 1] \subset \mathbb{R}^+$ according to (3).

$$S_{norm} = \frac{S_{raw}}{D_S}, \quad I_{norm} = \frac{I_{raw}}{D_I} \quad (3)$$

D_S, D_I denote the device-specific saturation counts. These normalized real-valued quantities are the inputs to the mapping function f .

Due to differences in grating technologies and filter characteristics, the discrete center wavelengths measured by the hyperspectral camera and the spectrometer are not perfectly aligned. We map between the wavelengths of the two devices with the following operations. Let

$$\Lambda_S = \Lambda_S^{(a)}_{a \in A} \text{ where each } \Lambda_S^{(a)} \in \mathbb{R}$$

$$\Lambda_I = \Lambda_I^{(b)}_{b \in B} \text{ where each } \Lambda_I^{(b)} \in \mathbb{R}$$

denote the sets of center wavelengths for the spectrometer and the camera, respectively, where $A = \{1, \dots, N_S\}$ and $B = \{1, \dots, N_I\}$ index the corresponding channels and N_S and N_I are the number of discrete channels for the spectrometer and hyperspectral camera, respectively. For each $b \in B$, we find the index of the spectrometer wavelength closest to the camera wavelength $\Lambda_I(b)$ and use it to define a mapping:

$$a^* = \arg \min_{a \in A} |\Lambda_S^{(a)} - \Lambda_I^{(b)}| \quad (4)$$

$$\Lambda_{cal}^{(b)} = \Lambda_S^{(a^*)} \quad (5)$$

which aligns each camera wavelength with its nearest neighbor in the spectrometer domain. (5) downsamples the spectrometer signal to the same number of wavelength values as the corresponding hyperspectral camera. This operation presupposes that the range of Λ_S spans the minimum and

maximum values of Λ_I in order to have good correspondences. In all further mentions of a spectral signal S_{raw} , we assume this signal is already subsampled to the dimension of Λ_{cal} , a subset of wavelengths corresponding between the two devices. Fig. 4a depicts this downsampling operation.

Next, we subtract off the device dark counts. These counts are acquired when the device is capped and no photons are able to enter the sensing aperture. Thus, the observed reading is an estimate of the intrinsic sensor noise. In practice, this value is found to be stable once the device reaches a steady-state operating temperature. Taking the minimum values at each wavelength yields an appropriate lower bound. Similarly, the dark counts are also normalized with respect to the digital counts. These values are then clamped to ≥ 0 .

$$S_{norm,z} = \max\left(0, S_{norm} - \frac{\min(S_{dark})}{D_S}\right) \quad (6)$$

$$I_{norm,z} = \max\left(0, I_{norm} - \frac{\min(I_{dark})}{D_I}\right) \quad (7)$$

To generate our approximation of $f : S \rightarrow I$, we expand the domain of this function to include the domain of the normalized spectra $S_{norm,z}$, and the location of the pixel on the photodetector array, r, c . As shown in [11], snapshot hyperspectral cameras tend to have non-negligible optical attenuation, especially near the edges of the image. For now, we will leave this function denoted by the generic $f(S_{norm,z}, r, c)$.

To construct the white reference for the camera, we utilize the following formulation:

$$I_{white}(r, c) = f(S_{norm,z}, r, c) \quad (8)$$

where $r \in \{0, \dots, H\}$ and $c \in \{0, \dots, W\}$

(8) yields a white reference cube in the shape of the target hyperspectral datacube derived from the spectrometer white reference spectra. We are now able to estimate the calibrated hyperspectral datacube I_{cal} , using (1), but generalized to hyperspectral datatypes.

$$I_{cal} = \frac{\frac{I_{raw}}{D_I} - \frac{I_{dark}}{D_I}}{\frac{I_{white}}{D_I} - \frac{I_{dark}}{D_I}} \quad (9)$$

In (9), the integration time from the calculation of the white reference is excluded. Considering that this data was collected with a fixed integration time t_{base} , we simply consider a scale factor given as the ratio of the base time to the current time. This is motivated under the assumption that the digital counts increase linearly with changes in the integration time.

$$t_{scale} = \frac{t_{base}}{t_{new}} \quad (10)$$

The scale factor from (10) is needed whenever the integration time of the spectrometer or hyperspectral camera differs from the original set integration time. Adding an integration scale factor yields (11), which will be used to recover the reflectance datacube. This equation is still dependent on the mapping function f , which we will learn an approximate solution to using empirical data in the following section. Graphically, this calibration process is shown in Fig. 3.

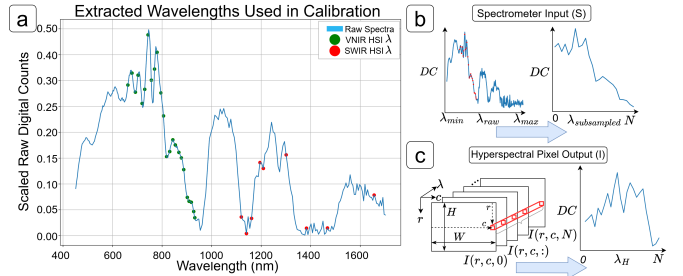


Fig. 4: (a) Single digital number normalized white reference measure from the HYPER DRIVE spectrometer system with the sampled spectra of the VNIR HSI (green dots) and SWIR HSI (red dots). (b) Downsampling of raw spectra to the wavelengths associated with the (c) temporally synchronized hyperspectral image.

$$I_{cal} = \frac{\left(\frac{I_{raw}}{D_I} \times t_{scale,I}\right) - \left(\frac{I_{dark}}{D_I} \times t_{scale,I}\right)}{f\left(\left(\frac{S_{raw}}{D_S} - \frac{S_{dark}}{D_S}\right) \times t_{scale,S}, \forall r, \forall c\right) - \left(\frac{I_{dark}}{D_I} \times t_{scale,I}\right)} \quad (11)$$

C. Modeling Spectral Mapping Functions

Given the observations of the correlation, or lack thereof, between spectra, we attempt to fit two different kinds of regression models on a per-pixel basis, using a) Multiple Linear Regression (MLR) and b) Multi-Layer Perceptron (MLP) Regression. We selected these two approaches since they constitute linear and non-linear modeling to the white reference reconstruction problem. We first begin by introducing the generic formulation of the problem. These formulations are applied to each wavelength pairing of the spectrometer and hyperspectral camera. Given its ability to handle spatial context, a more complex convolutional neural network (CNN) outwardly appears to be well-suited to this reconstruction task. Our experimentation with these CNN-based models showed a lack of convergence, likely due to requirements for a larger dataset to train a model of this complexity. Hence, we adhered to a per-pixel reconstruction of white reference data.

a) *Model Inputs & Outputs*: We define the input vector to this problem for a single sample as the reduced order spectral signal as calculated with (5). From a complexity perspective, downsampling also greatly helps save on computation time and model storage. Note that in the following equation, we only consider the output of a single pixel at a time. Our modeling efforts seek to find a model which associates a set of input spectra with Λ_I wavelength values, with the pixel response at location (r, c) that also contains Λ_I values. For the general case, we consider there are P available samples from the rooftop collection data. The problem is modeled graphically in Fig. 4. We also assume that all values of S_p and I_p are $\in \mathbb{R}^+$.

b) *Regression Modeling*: The first way of conceptualizing the relationship between input and output spectra is that of a multi-output regression model. Given the dataset collected during the rooftop system calibration campaign, we have a set of spectrometer measurements, each correlated with one of the white reference measurements taken for a given cube. MLR attempts to learn a weight matrix with ordinary least squares to map between the two sensors with minimal error.

c) *Multi-Layer Perceptron*: We also considered a deep learning model based on an MLP to approximate f . The general structure of these models followed an encoder-decoder architecture, where the encoder learns to encode an input signal

to a reduced-dimension latent space before reprojection to a higher dimension. The learned network consisted of 4 fully-connected layers with 24, 10, 10, and 24 neurons, respectively. A ReLU function is placed after each of the fully-connected layers.

When training these models, the loss term is crucial in accurately reconstructing the white reference hypercube. For instance, a signal that is reconstructed with similar absorbance troughs and peaks, but at lower magnitude is more desirable than one with a low mean-squared error (MSE) but misplaces the location of spectral features. The use of spectral angle mapper (SAM) is well supported by the literature as a means to calculate the divergence between two spectral signals [17]. We construct a composite loss function which simultaneously considers SAM loss and MSE loss. The equations for MSE (12) and SAM (13) specific to this task are presented below.

$$\mathcal{L}_{MSE}(r, c) = \frac{1}{P} \sum_{p=1}^P (I_p(r, c) - I'_p(r, c))^2 \quad (12)$$

$$\mathcal{L}_{SAM}(r, c) = \arccos \left(\frac{\sum_{p=0}^P I_p(r, c) \cdot I'_p(r, c)}{\sqrt{\left(\sum_{p=0}^P I_p(r, c)\right)^2} \sqrt{\left(\sum_{p=0}^P I'_p(r, c)\right)^2}} \right) \quad (13)$$

The range of $MSE \in \mathbb{R}_+$; $SAM \in [0, \pi]$. Together, the loss functions are combined with a weighting term α to weight their contributions to the total loss for the current sample. For training these models, α is set to 0.1. The combined loss function is given in (14).

$$\mathcal{L}_{total} = \mathcal{L}_{MSE} + \alpha \mathcal{L}_{SAM} \quad (14)$$

As the models are fitted on a per-pixel basis, they are trained in parallel. GPU acceleration is not needed for the MLP training as loading data onto the GPU for each model training process often takes longer than the convergence time of the model on the CPU. Models are implemented with the Scikit-Learn library [18] and trained for 1000 epochs, with early stopping if the loss function does not improve by 0.0001 units within 10 successive epochs. The Adam optimizer [19] is used to update the weights during training.

d) Data Augmentation & Processing: The spectral data was augmented in two ways to increase the generality of the learned models. First, we replicated the data by a factor of 3. We then applied scaling noise to both the input and output data drawn from a uniform distribution, adding up to $\pm 10\%$ scaling: $s_p = (1 + \mathcal{U}(-0.1, 1.0)) * s_p$, where \mathcal{U} is a uniform probability distribution. We clipped the data to ensure no negative values were introduced. Secondly, with an event probability of $\mathbb{P} = 0.10$, random indices in the range $\{0, \dots, |\Lambda_{cal}|\}$ of both the input and output were zeroed. This strategy was used as a means to encourage the model to learn generalized weights that handle different spectral distributions. This data was then split into 80/10/10% training, validation,

and testing percentages, respectively.

V. MODELING RESULTS

We assess the performance of MLP and MLR models using generalized SAM and MSE metrics as error measures rather than loss functions. Metrics such as mean absolute error (MAE) and MSE are calculated per pixel and averaged over all samples in P_{test} , where I is the target and I' represents model predictions. The VNIR MLR model exhibits the lowest errors near the center of the detector, with errors increasing radially outward, while the SWIR camera shows circularly radiating errors due to intrinsic system noise.

MLP generally overestimates white spectra intensity compared to MLR but captures spectral troughs and peaks more accurately, aiding appropriate scaling. Overestimation is acceptable if consistent, as underestimation can result in target pixel values exceeding the $[0, 1]$ normalization range, leading to excessive clipping and loss of spectral detail. For scaled spectrum inputs below 10% of the device's dynamic range, MLR outperforms MLP in accuracy. However, MLP achieves reasonably correct magnitudes with some noise. Optimal model performance requires sufficient signal, suggesting future work on dynamically adjusting integration times.

Summary statistics in Table III, computed with the held-out test set, show that MLR yields lower MSE and MAE, aligning with its optimization via MSE. However, MLR exhibits greater spectral angle divergence (higher SAM values). In contrast, MLP models generate spectra with the same shape (characteristic peaks and troughs) as the reference spectra, with trade-offs in inference time and model size.

Given limited computational resources, models must balance storage and runtime efficiency. MLR models, with matrix-based weight storage, scale quadratically with input-output dimensions but are smaller and faster than MLPs. Though none of the models operate in real-time, stable ambient solar spectra reduce the need for ad-hoc white reference recalibration. This stability allows robotic systems to allocate hardware resources more effectively to other tasks, as the solar spectrum does not change on a per-minute basis.

A. Color Marker Recovery

To evaluate the proposed method, we tested its ability to produce consistent measurements for reflectance targets of known intensity using a calibrated color checker (Checkr 24, Spyder). The setup, shown in Fig. 5, includes the color checker and a Spectralon tile for establishing illumination ground truth. The bottom row of the checker contains grayscale tiles of decreasing intensity. Fig. 5a highlights these tiles, and Fig. 5c presents the calibrated image. Spectralon exhibits uniform intensity across the spectrum, and the grayscale tiles correctly decrease in intensity, albeit with the 20% and 0% intensities closer than expected due to the camera's poor low-light sensitivity. Calibration mitigates variations seen in the

TABLE III: Summarized Reconstruction Error Statistics

Model Type	Sensor	\overline{MSE}	$\sigma(MSE)$	\overline{MAE}	$\sigma(MAE)$	\overline{SAM}	$\sigma(SAM)$	Model Size/Pixel (MB)	Inference (s)
MLR	VNIR	0.0001	0.0000	0.0041	0.0006	0.2380	0.0272	0.0016	3.41
	SWIR	0.00000	0.0000	0.0008	0.0001	0.3376	0.0248	0.0018	1.38
MLP	VNIR	0.0019	0.0007	0.0290	0.0062	0.1807	0.0496	0.0299	5.34
	SWIR	0.0029	0.0018	0.0324	0.0137	0.2477	0.0556	0.0294	2.26

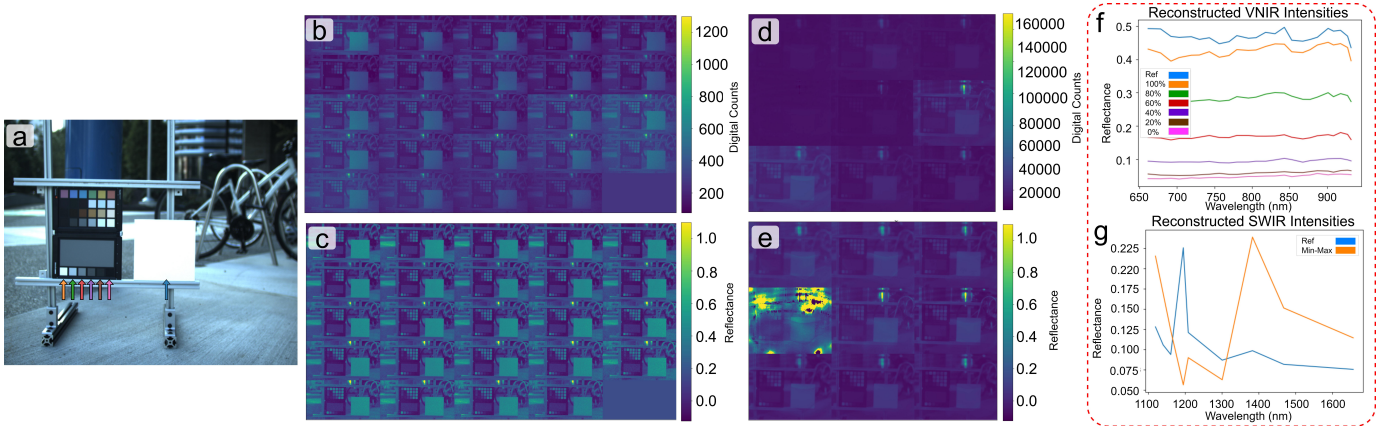


Fig. 5: (a) RGB image showing color checker (left) and Spectralon white reference target (right). The colored arrows correspond to the calibrated grayscale intensities seen in (f). (b) Raw demosaiced hyperspectral image with intensity in digital counts. Wavelengths increase from top-left to bottom right. (c) Reflectance calibrated image montage using the per-pixel MLP. Wavelengths increase from top-left to bottom right. (d) Raw digital count measurement from SWIR camera. (e) Reflectance calibrated image montage using the per-pixel MLP. Wavelengths increase from top-left to bottom right. (f) Reconstructed VNIR reflectance intensities for a 5-pixel neighborhood located in each target square. (g) Reconstructed SWIR reflectance intensities for Spectralon target in scene compared to naive min-max normalization.

uncalibrated datacube (Fig. 5b), stabilizing VNIR reference spectra to within 5% of the mean reflectance.

Similar trends are observed in SWIR calibration. Uncalibrated SWIR data (Fig. 5d) display significant intensity variation across wavelengths, while the calibration method (Fig. 5g) normalizes these discrepancies. However, low ambient illumination at 1195 nm results in a mean signal of $4.5 \pm 2\%$, creating calibration challenges due to low intensity and noise. Despite this, the calibration normalizes the white reference to within $\pm 3\%$, except for a spike at 1195 nm due to the aforementioned absorption in this band. This aligns with expectations for uniform intensity recovery.

In both VNIR and SWIR reconstructions, spectral values are below the ideal value of 1 for the Spectralon tile. Variations in average spectral intensity arise from factors such as atmospheric water content or shadows, as noted in prior research [20], [21]. Per-wavelength normalization remains consistent, as evidenced in Fig. 5. Techniques such as SAM, commonly used in remote sensing and machine learning [22], are robust against albedo variations, making them suitable for scenes with varying intensities.

VI. TERRAIN PARAMETER ESTIMATION

With the properly calibrated images, we consider the estimation of two terrain properties: vegetative health (VH) and soil moisture content (SMC) from a moving autonomous vehicle. For these experiments, the HYPER DRIVE system was mounted on a large field robot (Warthog, Clearpath Robotics) to collect outdoor imagery. When mounted, the camera array is 1.0 meter off the ground. Because the spectrometer samples a uniform white reference tile mounted rigidly to the platform, transient shocks or vibrations from UGV motion do not affect the calibration.

A. Vegetative Health Analysis

As a field robot is moving through an environment, understanding where vegetation is present can aid in gauging the feasibility of a selected path. For instance, if the target path is primarily composed of gravel or sand, vegetation can be

considered a non-viable area to traverse as vegetation is not normally present in cleared paths.

Normalized difference vegetation index (NDVI) is a well-studied and proven metric towards understanding the location of vegetation [23]. NDVI is a normalized comparison between two spectral reflectance bands with a range of $[-1, 1]$. Dense, healthy vegetation will yield a value near 1, while unhealthy vegetation will be less than 0.5. Values less than 0 suggest the presence of standing water.

Using our camera system, the NDVI is calculated along the channel axis with (15). The subscripts to I represent the wavelength bands extracted for the calculation of the index.

$$NDVI = \frac{I_{901} - I_{661}}{I_{901} + I_{661}} \quad (15)$$

With the NDVI values, we extract a statistical decision boundary between vegetation and non-vegetation using Otsu's method [24]. This method assumes there are exactly two classes and seeks to minimize intra-class variance for the binary decision problem. We prefer the binary Otsu over a multi-Otsu thresholding [25] implementation as we presume vegetation vs. non-vegetation is a useful differentiation for off-road navigation and has precedence in the literature [11]. No post-processing or contrast enhancements are applied to the NDVI scores.

The first column of Fig. 6 shows the raw RGB image for the scene. The next two columns show the calculated NDVI for both the raw and calibrated images. The final two columns show the Otsu segmentation performed on the calibrated and uncalibrated images, respectively. In the Otsu images, yellow indicates the pixels clustered as vegetation, while dark purple indicates non-vegetation.

Row (a) presents an RGB color checker scene with a background of vegetation and predominantly man-made structures. NDVI calculations reveal a strong contrast between vegetated areas and other regions. Reflective surfaces on the color checker edges produce non-zero NDVI readings. Otsu's method better segments vegetation compared to uncalibrated data, which has more pixels misclassified as vegetation. Rows (b) and (c), featuring dense undergrowth and shadows, exhibit lower raw HSI datacube intensity. Calibrated NDVI effectively handles

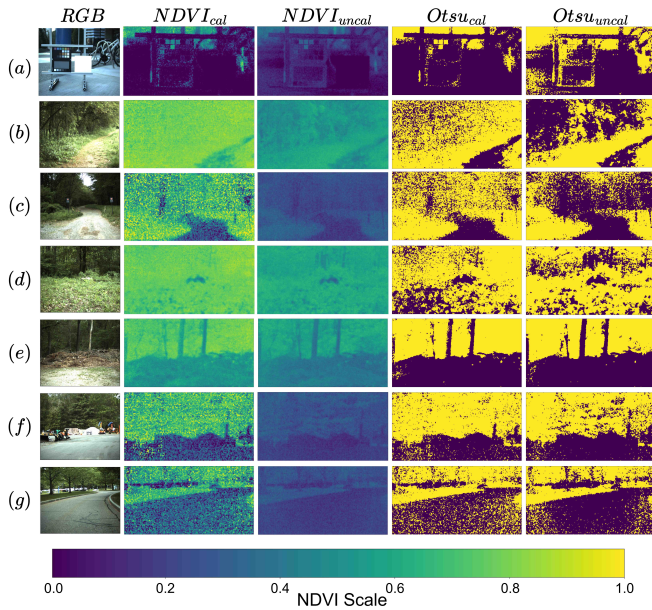


Fig. 6: NDVI values calculated using both the raw images and the proposed calibration method. These images are sampled from the HYPER DRIVE [16] dataset.

illumination variations, though darker areas in uncalibrated data are misclassified. Some speckling in calibrated NDVI is noted due to low signal strength.

Rows (d) and (e) show varying vegetative health, with living plants and dry leaves. Calibrated Otsu improves detection of verdant vegetation and anomalies like boulders (d). It also better identifies background trees as vegetation (e), compared to uncalibrated data. Rows (f) and (g) feature mixed vegetation and man-made structures. Calibration stabilizes vegetation NDVI in low-light scenarios (c), (f), and (g), while uncalibrated NDVI values fluctuate with illumination.

Otsu’s method provides insight into class separability but assumes only two distinct classes. In scenes with multiple classes, such as stressed vegetation or standing water, its heuristic nature may skew results. Nonetheless, our results show calibrated NDVI is a useful prior or input for semantic segmentation algorithms, enabling effective terrain analysis.

B. Soil Moisture Detection

Terrain moisture is also a significant area of interest to enable decision-making grounded in physical quantities. Towards understanding soil moisture content (SMC), we constructed a series of testbeds where soil moisture could be precisely controlled and measured in situ with probes. Fig. 7a shows the test setup in a gridded format of increasing soil moisture with HYPER DRIVE imaging the samples.

The testbed was constructed in a 3×3 grid, with separations between adjacent soil samples to control the water content and minimize flow between cells. One box was kept dry while the remaining boxes had increasing amounts of water. Ground truth for SMC was acquired by measuring soil relative humidity (RH) (%) using a surface-penetrating probe at 10 points in each cell. Measurements ranged from 0% - 48.1% across the testbed. The testbed was imaged three times as the robot was teleoperated up to the edge of the boxes.

After acquiring the ground truth measurements, we masked the images and extracted the pixels contained within each box.

Fig. 7b shows the representative spectral signature for each of the samples, plotted with \pm one standard deviation from the mean box spectra. Qualitatively, the VNIR spectra show only changes in magnitude with the increasing soil moisture content. Wetter sand has a marginal albedo change. However, it is important to recall that changes in magnitude are not stable predictors, especially in partially shadowed environments. The SWIR shows more promising features, particularly the change at ≈ 1300 nm.

Inspired by the formulation of NDVI, we formulate an optimization problem to quantitatively choose a band pairing for a SMC index using the wavelength of the VIS-SWIR HSI system. We hypothesize that maximizing the distance between the maximum and minimum index value across bands will best allow for differentiation of intermediary moisture contents. In (16), the mathematical optimization of the band selection problem is presented. We optimize the spectral index by considering the pair of band combinations yielding the maximum difference between the dry and wet sand.

$$f(i, j) = \left\| \frac{p_{i,wet} - p_{j,wet}}{p_{i,wet} + p_{j,wet}} - \frac{p_{i,dry} - p_{j,dry}}{p_{i,dry} + p_{j,dry}} \right\|_2 \quad (16)$$

$$\underset{i, j}{\text{maximize}} f(i, j)$$

$$\text{subject to } i \in \Lambda, j \in \Lambda, i \neq j.$$

$$SMC = \frac{I_{1300} - I_{1119}}{I_{1300} + I_{1119}} \quad (17)$$

p is the set of pixels, with values at wavelength i or j . The wet and dry pixels are extracted from boxes 1 and 9, respectively. Λ is the set of wavelengths in the registered datacube. The optimized wavelength values align with the spectral absorption features present between ≈ 1150 nm and ≈ 1350 nm in the solar absorption spectra.

Extracting the optimal wavelength values yields (17), which calculates the normalized difference between the two wavelength channels. The value of SMC here is dimensionless. In order to convert the SMC to an actual physical quantity, we consider the SMC index as the domain of a linear regression function. Fig. 7d depicts the results of this fit model.

From Fig. 7, the model shows a strong ability to distinguish the dry box from the saturated box. Predictions of the intermediary values are distinguishable, but not to the same degree as the extreme situations of dry and saturated. The distribution of the spectral values in each of these boxes shows an overlapping range of values. Given the results, it may be worth considering soil moisture as a three-class problem: dry, moist, and saturated. Considering these classes, vehicles could tailor their risk tolerances depending on which moisture levels constitute go or no-go decisions through a physically informed terramechanics model.

One of the foremost areas of future work is the extension of these methods to other soil types. The results indicate a strong ability to discern surface properties beyond broad semantic labels such as *sand* or *soil*. These results are produced on the assumption that the terrain class is already known. With this knowledge, SMC could then be applied to specific regions where the moisture likely affects terramechanics such as sand, soil, and clay. The prior knowledge of terrain class also reduces the computational load on the characterization algorithm, as only certain regions will need this property calculated. Additionally, per-class property modeling allows

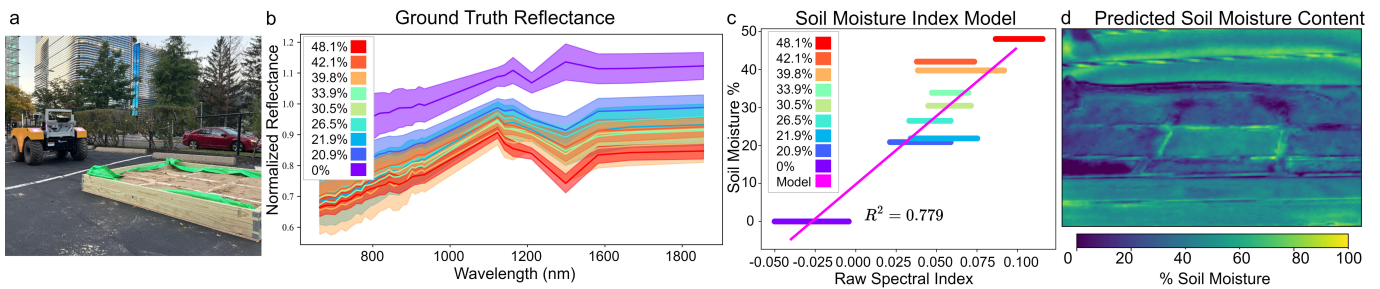


Fig. 7: (a) HYPERSPECTRAL DRIVE mounted to the field robot scanning sample testbed. (b) Ground truth spectral signatures for the calibration testbeds. (c) Soil moisture content in each testbed cell with fit regression line. (d) Predicted soil moisture using the regression model.

for different band combinations to be optimized for different soil types. While 1300 and 1119 nm are optimal for sand, the presence of organic material in other surface types might necessitate the integration of bands from the VNIR domain.

Although these example classification processes use two bands each, they leverage different parts of the spectrum, and show how the same calibrated spectral data can be utilized for both terrain identification and parameterization. Calculating these spectral indices takes 30 ms on a CPU per datacube. This speed (30 frames per second) supports real-time processing. This ability to derive a real, physical quantity from the environment with minimal labels, minimum compute requirements, and simple modeling strongly motivates further investigation into HSI in field robotics.

VII. CONCLUSIONS

We presented a system architecture for acquiring and calibrating hyperspectral images from a moving robotic platform. The proposed method uses spectrometer-hyperspectral joint calibration to mitigate the effects of varying illumination and enable consistent reflectance measurements without calibration targets. Using this calibration, we demonstrated the estimation of two terrain-relevant properties: vegetative health and soil moisture content.

This work provides a practical and computationally efficient framework for integrating hyperspectral sensing into field robotics. Future efforts will focus on extending these methods to a broader range of terrain types and developing autonomous strategies for identifying and responding to environmental conditions that affect traversability.

REFERENCES

- [1] F. Islam, M. Nabi, and J. E. Ball, "Off-road detection analysis for autonomous ground vehicles: A review," *Sensors*, vol. 22, no. 21, p. 8463, 2022.
- [2] R. Schmid, D. Atha, F. Schöller, S. Dey, S. Fakoorian, K. Otsu, B. Ridge, M. Bjelonic, L. Wellhausen, M. Hutter *et al.*, "Self-supervised traversability prediction by learning to reconstruct safe terrain," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 12 419–12 425.
- [3] N. Hanson, T. Kelestemur, J. Berman, D. Ritzenhoff, and T. Padir, "Hyperbot—a benchmarking testbed for acquisition of robot-centric hyperspectral scene and in-hand object data," in *2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*. IEEE, 2022, pp. 1–5.
- [4] D. Manolakis, E. Truslow, M. Pieper, T. Cooley, and M. Brueggeman, "Detection algorithms in hyperspectral imaging systems: An overview of practical algorithms," *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 24–33, 2013.
- [5] S. Matteoli, M. Diani, and G. Corsini, "A tutorial overview of anomaly detection in hyperspectral images," *IEEE Aerospace and Electronic Systems Magazine*, vol. 25, no. 7, pp. 5–28, 2010.
- [6] S. Nann and C. Riordan, "Solar spectral irradiance under clear and cloudy skies: Measurements and a semiempirical model," *Journal of Applied Meteorology and Climatology*, vol. 30, no. 4, pp. 447–462, 1991.
- [7] T. Yamanouchi and M. Tanaka, "Absorption properties of the near-infrared water vapor bands," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 34, no. 6, pp. 463–472, 1985.
- [8] T. W. Ray, "A faq on vegetation in remote sensing," *California: Div. of Geological and Planetary Sciences California Institute of Technology*, 1994.
- [9] T. Hakala, E. Honkavaara, H. Saari, J. Mäkynen, J. Kaivosoja, L. Pesonen, and I. Pölonen, "Spectral imaging from UAVs under varying illumination conditions," in *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. International Society for Photogrammetry and Remote Sensing (ISPRS), 2013.
- [10] A. Wendel and J. Underwood, "Illumination compensation in ground based hyperspectral imaging," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 129, pp. 162–178, 2017.
- [11] R. Khanna, I. Sa, J. Nieto, and R. Siegwart, "On field radiometric calibration for multispectral cameras," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 6503–6509.
- [12] G. Blackburn and F. Vignola, "Spectral distributions of diffuse and global irradiance for clear and cloudy periods," in *World Renewable Energy Forum, Denver, CO*, 2012.
- [13] P. Geladi, J. Burger, and T. Lestander, "Hyperspectral imaging: calibration problems and solutions," *Chemometrics and intelligent laboratory systems*, vol. 72, no. 2, pp. 209–217, 2004.
- [14] P. Horstrand, R. Guerra, A. Rodríguez, M. Díaz, S. López, and J. F. López, "A UAV platform based on a hyperspectral sensor for image capturing and on-board processing," *IEEE Access*, vol. 7, pp. 66 919–66 938, 2019.
- [15] N. Hanson, M. Shaham, D. Erdoğmuş, and T. Padir, "Vast: Visual and spectral terrain classification in unstructured multi-class environments," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3956–3963.
- [16] N. Hanson, S. Hibbard, B. Pyatski, C. DiMarzio, and T. Padir, "Hyperdrive: Visible-short wave infrared hyperspectral imaging datasets for robots in unstructured environments," in *2023 13th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*. IEEE, 2023.
- [17] F. A. Kruse, A. Lefkoff, J. Boardman, K. Heidebrecht, A. Shapiro, P. Barloon, and A. Goetz, "The spectral image processing system (sips)—interactive visualization and analysis of imaging spectrometer data," *Remote sensing of environment*, vol. 44, no. 2-3, pp. 145–163, 1993.
- [18] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [20] C. J. Tucker, "Remote sensing of leaf water content in the near infrared," *Remote sensing of Environment*, vol. 10, no. 1, pp. 23–32, 1980.
- [21] G. A. Carter, "Primary and secondary effects of water content on the spectral reflectance of leaves," *American journal of botany*, vol. 78, no. 7, pp. 916–924, 1991.
- [22] S. Khanna, A. Palacios-Orueta, M. L. Whiting, S. L. Ustin, D. Riaño, and J. Litago, "Development of angle indexes for soil moisture estimation, dry matter detection and land-cover discrimination," *Remote sensing of environment*, vol. 109, no. 2, pp. 154–165, 2007.
- [23] J. W. Rouse, R. H. Haas, J. A. Schell, D. W. Deering *et al.*, "Monitoring vegetation systems in the great plains with erts," *NASA Spec. Publ.*, vol. 351, no. 1, p. 309, 1974.
- [24] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE transactions on systems, man, and cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.
- [25] P.-S. Liao, T.-S. Chen, P.-C. Chung *et al.*, "A fast algorithm for multilevel thresholding," *J. Inf. Sci. Eng.*, vol. 17, no. 5, pp. 713–727, 2001.