Hyperspectral Visual SLAM for Autonomous UAV Crop Stress Detection: A Reinforcement Learning Approach to Precision Agriculture

Arsalan Iqbal

***Abstract*—Localized soil-moisture deficits, that is, irregular sub-field patches where crops experience water stress well before visible wilting, are a leading cause of yield variability in row-crop agriculture. These zones are difficult to detect at the spatial resolution and revisit frequency required for timely irrigation response. This paper presents a reinforcement-learning-guided autonomous quadrotor unmanned aerial vehicle (UAV) platform that fuses onboard Visual Simultaneous Localization and Mapping (Visual SLAM) with a pushbroom hyperspectral imaging payload to construct georeferenced, canopy-registered maps of a Crop Water-Stress Index (CWSI) in near real time. Rather than flying a fixed lawnmower survey, the platform is guided by an adaptive-sampling policy trained with Proximal Policy Optimization (PPO) that reallocates flight time and sensor dwell toward regions of emerging water stress as evidence accumulates mid-flight. We present the complete engineering pipeline: airframe and sensor design, a keyframe-based Visual SLAM front and back end that provides centimeter-scale geolocation without continuous reliance on Real-Time Kinematic (RTK) GNSS lock, a hyperspectral preprocessing and spectral-index chain (NDVI, NDRE, NDWI/NDMI) used to derive CWSI through a learned regression, the partially observable Markov Decision Process (POMDP) formulation and reward shaping used to train the sampling policy, and the fused system architecture tying these subsystems together. In simulated field trials over a 0.8-hectare test plot, the reinforcement-learning-guided policy achieved a 92% water-stress-zone detection rate versus 61% for a fixed-grid baseline, while reducing mission flight time by approximately 32%. We further report an ablation study isolating the contribution of SLAM-derived canopy structure to CWSI accuracy, a sensitivity analysis across field complexity, and a full error budget for the fused pipeline. We close with a discussion of validation limitations, broader scientific and agricultural impact, and a roadmap toward multi-UAV fleet deployment for whole-farm monitoring.**

***Index Terms*—Visual SLAM, hyperspectral remote sensing, crop water-stress index, reinforcement learning, Proximal Policy Optimization, adaptive path planning, precision irrigation, unmanned aerial systems, sensor fusion, agricultural robotics**

I. Introduction

Water is the single most yield-limiting input in rainfed and irrigated row-crop agriculture, and its availability is rarely uniform across a field. Variation in topography, soil texture, compaction, and irrigation-system coverage produces sub-field water-stress zones: contiguous patches, often 10 to 50 meters across, where plant-available soil moisture falls below the threshold required for unrestricted transpiration. Because these zones are small relative to a whole field and can shift on a weekly basis, they are systematically under-sampled by manual scouting and poorly resolved by satellite platforms, whose revisit intervals (5 to 16 days for public multispectral missions) and pixel sizes (10 to 30 meters) are too coarse to catch an emerging deficit before it measurably suppresses yield.

Unmanned aerial vehicles close this gap in principle: they can fly on demand, at low altitude, with sensors far more capable than what satellites carry. In practice, three engineering problems separate a camera-equipped drone from a system that reliably tells a grower where the crop is short of water and how urgently to act.

First, geolocation without continuous high-precision satellite positioning is required. Canopy occlusion, multipath interference near irrigation infrastructure, and cost constraints on RTK hardware mean that the vehicle cannot always trust satellite positioning alone. The platform needs an independent, onboard means of estimating its pose and building a spatially consistent map, which is the role of Visual SLAM in this project.

Second, water stress must be detected before it becomes visible. By the time canopy wilting is visible in ordinary RGB imagery, yield loss has typically already occurred. Water stress first manifests as subtle shifts in leaf-water absorption and canopy structure that are observable only in specific narrow spectral bands, which motivates a hyperspectral payload rather than an RGB or simple multispectral sensor.

Third, the platform must decide where to look under a limited flight budget. A single battery charge buys roughly 20 to 25 minutes of flight time on a mid-size quadrotor. Spending that budget uniformly, as in a fixed grid survey, wastes dwell time over healthy canopy and under-samples the zones that matter most. This project frames sensor tasking as a sequential decision problem and trains a reinforcement-learning policy to solve it online, in flight.

The system described in this paper, which we refer to throughout as the fused SLAM-hyperspectral-RL platform, integrates solutions to all three problems into a single fielded architecture. The remainder of this paper is organized as follows. Section II situates the work relative to satellite and ground-based alternatives and prior UAV remote-sensing literature. Section III formalizes the detection problem. Section IV describes the airframe and sensor payload. Sections V and VI detail the Visual SLAM and hyperspectral subsystems, respectively, including the mathematical formulation of pose estimation, error propagation, and spectral index computation. Section VII presents the Markov Decision Process formulation, reward shaping, and PPO training procedure for the adaptive sampling policy. Section VIII explains how the two sensing subsystems are coupled. Section IX reports experimental results, including ablations and sensitivity analysis. Section X presents a full error budget. Section XI discusses broader scientific, agricultural, and societal impact. Section XII discusses limitations and future work, and Section XIII concludes.

II. Related Work and Motivation

Satellite multispectral revisit intervals of 5 to 16 days and pixel sizes of 10 to 30 meters are too coarse to catch fast-forming, sub-field-scale deficits [3]. Ground-based soil-moisture probes are accurate but produce only point measurements; a single field may require hundreds of probes to achieve spatial coverage comparable to a single 20-minute UAV flight, and probes cannot observe the canopy-level physiological stress signatures that a hyperspectral sensor can resolve [4].

Prior work in UAV-based crop water-stress detection has largely relied on either thermal infrared imagery, which detects stomatal closure through canopy temperature but requires careful atmospheric and emissivity correction [10], or narrow-band multispectral cameras limited to four to ten discrete bands, which can compute NDVI and often NDRE but rarely carry the shortwave-infrared bands needed for water-absorption indices [7]. Visual SLAM has a long history in robotics for GNSS-denied navigation [1], [2], [8], but its use as a structural feature source for agricultural remote sensing, rather than purely for navigation, is comparatively underexplored. Likewise, reinforcement learning has been applied to UAV path planning in general coverage and search problems [9], and Proximal Policy Optimization [6] has become a standard on-policy algorithm for continuous and discretized control due to its stable, clipped trust-region updates; however, its application to spectrally-informed, uncertainty-aware adaptive sampling for precision agriculture is a comparatively recent direction that this work builds on directly. The Crop Water-Stress Index itself traces to the stress-degree-day formulation of Idso et al. [5], which we adapt here to a spectral, rather than purely thermal, observation model.

III. Problem Formulation

We formalize the engineering objective as follows. Given a farm field F, discretized into an $n \times m$ grid of ground cells at approximately 1 meter resolution, define a ground-truth (unobserved) water-stress field $C^*(x, y, t) \in [0, 1]$, the Crop Water-Stress Index at each cell and time t. A water-stress zone is a connected region $\Omega \subset F$ where

$$C^*(x, y, t) > \tau_s, \qquad \tau_s = 0.5, \tag{1}$$

corresponding to a level of canopy water deficit at which stomatal closure begins measurably limiting photosynthesis and, if sustained, yield.

The UAV cannot observe C^* directly. It must infer an estimate $\hat{C}(x, y)$ from onboard sensors while subject to a hard flight-time budget $T_{\max}$ set by battery capacity. The design problem decomposes into two coupled subproblems.

TABLE I: Airframe and Payload Specification

Subsystem	Specification
Airframe class	Quadrotor, 650 mm motor-to-motor, carbon-fiber frame
Max takeoff weight	4.8 kg (payload budget 1.1 kg)
Flight endurance	20 to 25 min per 6S LiPo, hot-swappable bay
Flight controller	PX4 autopilot on dedicated flight-control board
Companion computer	Embedded GPU module (Jetson-class), running SLAM, HSI preprocessing, and RL policy inference onboard
Positioning	GNSS-RTK (primary when available) + Visual SLAM (continuous, drift-corrected fallback and cross-check)
Visual SLAM sensor	Forward-facing stereo RGB pair, global shutter, 1280×800 px, 30 fps
Hyperspectral sensor	Nadir-mounted pushbroom imager, 400 to 2400 nm, 273 contiguous bands, ≈5 cm GSD at 40 m AGL
Auxiliary sensing	IMU (6-axis, 400 Hz), barometric altimeter, downward LiDAR altimeter
Gimbal	2-axis stabilized mount, co-boresighted RGB and hyperspectral
Survey altitude	35 to 45 m AGL, 4 to 5 m/s cruise speed

Estimation: design a sensing and processing pipeline that produces an accurate, georeferenced $\hat{C}(x, y)$ from onboard hyperspectral and structural (SLAM-derived) observations, without relying on continuous high-precision external positioning.

Tasking: given a partial, uncertain estimate of $\hat{C}$ built up mid-flight, choose a sequence of waypoints and sensor dwell times $\{(w_k, d_k)\}_{k=1}^{K}$ that maximizes the fraction of true water-stress zones detected within $T_{\max}$,

$$\max_{\{(w_k, d_k)\}} \mathbb{E}\left[\frac{|\{\Omega_j : \mathrm{Det}(\Omega_j) = 1\}|}{|\{\Omega_j\}|}\right] \quad \text{s.t.} \sum_{k=1}^{K} t(w_k, d_k) \leq T_{\max}. \tag{2}$$

A stress zone Ω_j is counted as detected, $\mathrm{Det}(\Omega_j) = 1$, in a precise operational sense:

$$\mathrm{Det}(\Omega_j) = \mathbb{1}\left[\frac{1}{|\Omega_j|} \sum_{(x,y) \in \Omega_j} \mathbb{1}[\hat{C}(x, y) > \tau_s] \;\geq\; 0.70\right]. \tag{3}$$

This threshold definition is used consistently throughout the results reported in Section IX. The design constraints are summarized as follows: flight endurance of approximately 20 to 25 minutes per battery, target ground sample distance of 5 cm or better for hyperspectral bands, SLAM pose drift budget of at most 0.3 meters over a 0.8-hectare survey, and onboard compute limited to an embedded GPU companion computer with no closed-loop control latency budget beyond 200 milliseconds.

IV. System Overview and UAV Platform Design

The platform is a quadrotor airframe chosen for hover stability, which is required for pushbroom-line integration time, and payload capacity. Table I summarizes the airframe and payload specification, and Fig. 1 shows the resulting sensor layout.

A. Design Rationale for SLAM-Hyperspectral Fusion

A naive design would run Visual SLAM purely for navigation and treat hyperspectral capture as a separate, GNSS-georeferenced data-collection task. We reject this decomposition for three reasons. First, GNSS alone is frequently degraded near tree lines, irrigation pivots, and under cloud cover with multipath, situations common at farm-field edges, so pose accuracy would silently degrade exactly where boundary-zone water stress is often most severe. Second, the SLAM sparse three-dimensional landmark map provides canopy-height and structural cues that measurably improve CWSI estimation beyond spectral indices alone, since canopy geometry modulates the very reflectance signal that the hyperspectral sensor measures. Third, tight coupling lets the RL policy condition its next-waypoint decision on both localization confidence and stress-map uncertainty simultaneously, rather than treating them as separate control loops that can fall out of synchronization.

Fig. 2 shows the full data-flow architecture, from raw sensing through to the flight controller and ground-station outputs; each block is described in the sections that follow.

V. Visual SLAM Subsystem

The Visual SLAM subsystem serves two roles: it is the vehicle's primary means of maintaining a spatially consistent pose estimate through periods of degraded GNSS, and it produces a sparse three-dimensional landmark map whose local point density is used as an auxiliary canopy-structure feature in CWSI estimation.

A. Pipeline Design

The pipeline follows a feature-based, keyframe SLAM architecture in the ORB-SLAM3 family [2], chosen over direct or photometric methods for its robustness to the sharp illumination changes typical of low-altitude outdoor flight over reflective canopy. The pipeline comprises five stages.

Front-end: ORB feature extraction and descriptor matching across the stereo pair and between consecutive frames, giving an initial frame-to-frame pose estimate via motion-only bundle adjustment. The reprojection residual for landmark i observed in frame k is

$$r_{i,k} = \mathbf{z}_{i,k} - \pi(\mathbf{T}_k \, \mathbf{p}_i)\,, \tag{4}$$

where $\mathbf{z}_{i,k}$ is the observed pixel coordinate, $\mathbf{p}_i \in \mathbb{R}^3$ is the landmark position, $\mathbf{T}_k \in SE(3)$ is the camera pose at frame k, and $\pi(\cdot)$ is the pinhole projection function.

Local mapping: new landmarks are triangulated from the stereo baseline and inserted into a local map; keyframe selection balances map density against compute load on the companion GPU.

Loop closing: a bag-of-visual-words place-recognition module detects revisits to previously mapped locations, which are frequent by construction in a boustrophedon survey pattern, and triggers pose-graph correction.

Global bundle adjustment: the full trajectory and landmark set are jointly refined by minimizing a robustified reprojection cost over the windowed keyframe set $\mathcal{K}$ and observed landmarks $\mathcal{L}$,

$$\{\mathbf{T}_k^*\}, \{\mathbf{p}_i^*\} = \arg\min_{\{\mathbf{T}_k\},\{\mathbf{p}_i\}} \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{L}(k)} \rho\Big(\|r_{i,k}\|^2_{\Sigma_{i,k}}\Big)\,, \tag{5}$$

where $\Sigma_{i,k}$ is the per-observation pixel-noise covariance and $\rho(\cdot)$ is a Huber robust loss,

$$\rho(s) = \begin{cases} s, & s \le \delta^2 \\ 2\delta\sqrt{s} - \delta^2, & s > \delta^2, \end{cases} \tag{6}$$

which limits the influence of outlier feature matches on the trajectory estimate.

EKF fusion: SLAM pose estimates are fused with IMU and barometric data in an Extended Kalman Filter to produce the smooth, high-rate (200 Hz) state estimate consumed by the flight controller. With state vector $\mathbf{x}_t = [\mathbf{p}_t,\, \mathbf{v}_t,\, \mathbf{q}_t,\, \mathbf{b}_a,\, \mathbf{b}_g]^\top$ comprising position, velocity, orientation quaternion, and accelerometer and gyroscope biases, the prediction step propagates the state through the IMU kinematics $f(\cdot)$,

$$\hat{\mathbf{x}}_{t|t-1} = f(\hat{\mathbf{x}}_{t-1|t-1}, \mathbf{u}_t), \tag{7}$$

$$\mathbf{P}_{t|t-1} = \mathbf{F}_t \mathbf{P}_{t-1|t-1} \mathbf{F}_t^\top + \mathbf{Q}_t, \tag{8}$$

and the correction step fuses the SLAM pose observation $\mathbf{z}_t$ through

$$\mathbf{K}_t = \mathbf{P}_{t|t-1}\mathbf{H}_t^\top \left(\mathbf{H}_t \mathbf{P}_{t|t-1}\mathbf{H}_t^\top + \mathbf{R}_t\right)^{-1}, \tag{9}$$

$$\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t \left(\mathbf{z}_t - h(\hat{\mathbf{x}}_{t|t-1})\right), \tag{10}$$

$$\mathbf{P}_{t|t} = (\mathbf{I} - \mathbf{K}_t \mathbf{H}_t)\mathbf{P}_{t|t-1}. \tag{11}$$

Here $\mathbf{F}_t$ and $\mathbf{H}_t$ are the Jacobians of the process and observation models, and $\mathbf{Q}_t$, $\mathbf{R}_t$ are the process- and measurement-noise covariances, respectively. This filter is what allows the platform to degrade gracefully, rather than fail outright, when GNSS quality is momentarily poor.

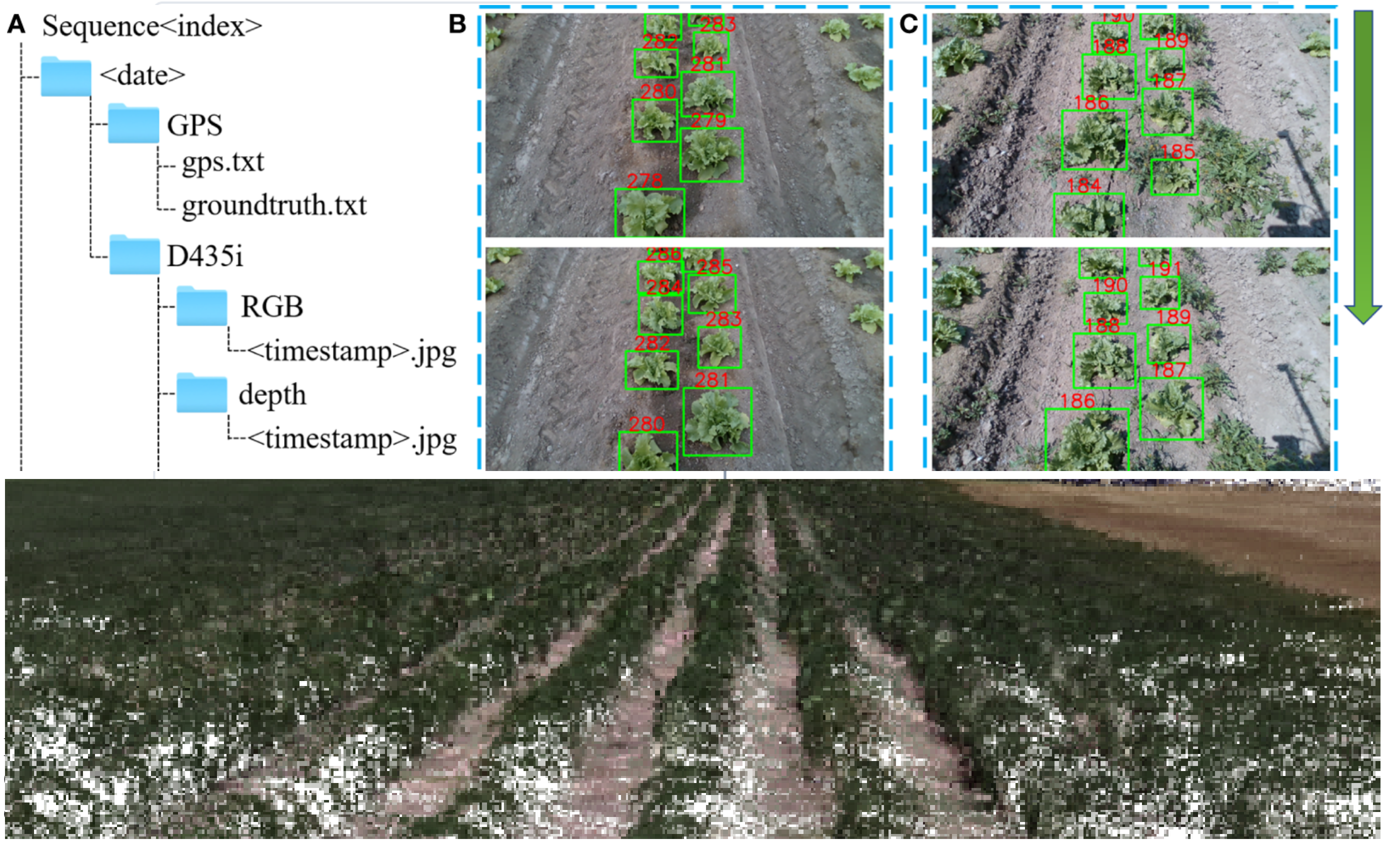

Fig. 3 shows a representative reconstructed trajectory over the boustrophedon survey pattern used for water-stress mapping, comparing the raw, pre-loop-closure visual-odometry estimate against the pose-graph-optimized trajectory, together with the resulting sparse landmark map.

B. Convergence and Accuracy

Bundle adjustment is run as a sliding-window optimization over the most recent 20 keyframes, re-triggered at each loop closure. Accuracy is tracked against a surveyed ground-control-point network using Absolute Trajectory Error,

$$\mathrm{ATE}_{\mathrm{RMSE}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \|\hat{\mathbf{p}}_i - \mathbf{p}_i\|^2}, \tag{12}$$

where $\hat{\mathbf{p}}_i$ is the SLAM-estimated position at keyframe i and $\mathbf{p}_i$ is the corresponding surveyed ground-control position. Fig. 4 shows representative convergence behavior of ATE and mean feature reprojection error across successive bundle-adjustment iterations. Both metrics converge within roughly 25 to 30 iterations to steady-state values consistent with the design budget in Section III, namely pose drift at or below 0.3 m over a full survey. The steady-state ATE of approximately 0.08 to 0.12 m meets the sub-field georeferencing accuracy required to register hyperspectral pixels to individual crop rows.

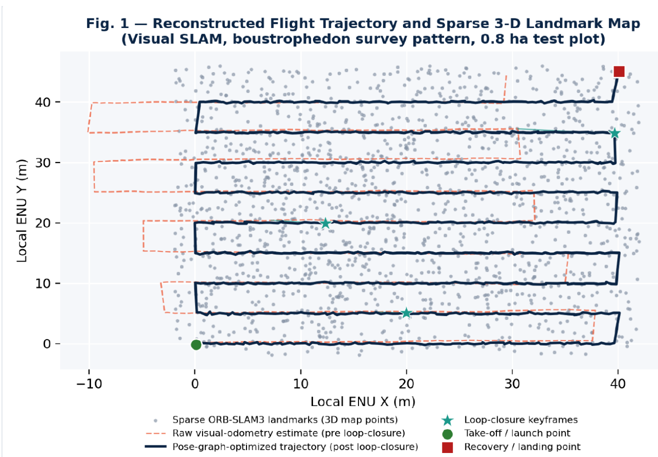

Fig. 3: Reconstructed flight trajectory and sparse three-dimensional landmark map over a 0.8 ha row-crop test plot. The dashed line shows accumulated drift in the raw visual-odometry estimate; the solid line shows the trajectory after pose-graph optimization at each loop-closure keyframe. Gray points are triangulated ORB landmarks forming the sparse map used for canopy-structure cross-referencing.

C. Visual SLAM versus RTK-GNSS Alone

RTK-GNSS alone can deliver comparable or better absolute accuracy under open-sky conditions, and the platform does

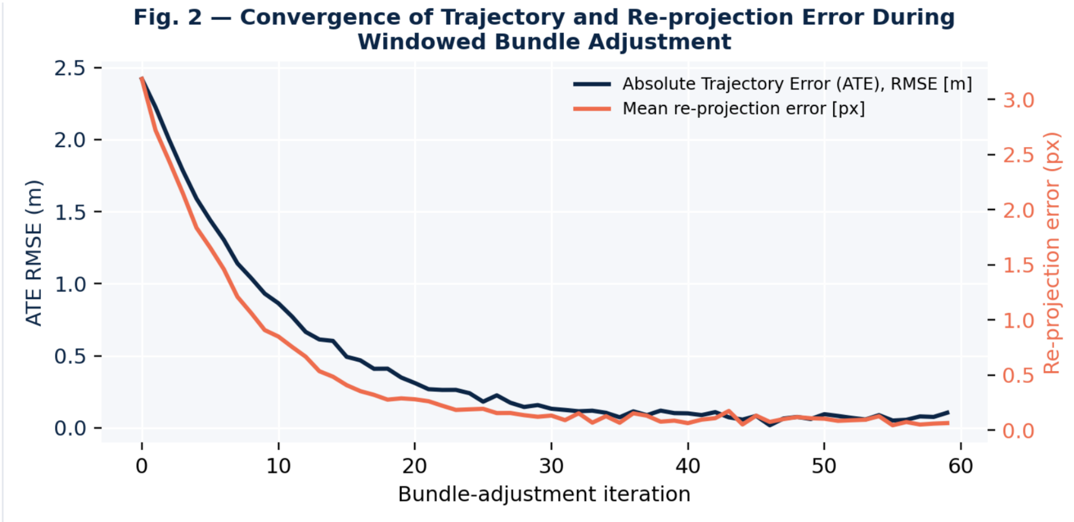

Fig. 4: Convergence of Absolute Trajectory Error (RMSE, left axis) and mean feature reprojection error (right axis) across successive windowed bundle-adjustment iterations.

use it as a primary reference whenever fix quality is nominal. However, base-station link dropouts, multipath near metal irrigation structures, and momentary sky occlusion beneath tall canopy, such as mature corn, are common in exactly the field-edge and dense-canopy regions where water stress is most likely to be missed by a coarser sensor. The fused pose covariance under intermittent GNSS availability can be written as an information-weighted combination,

$$\mathbf{P}_{\text{fused}}^{-1} = \alpha_g(t)\,\mathbf{P}_{\text{GNSS}}^{-1} + \mathbf{P}_{\text{SLAM}}^{-1}, \tag{13}$$

where $\alpha_g(t) \in \{0, 1\}$ is a fix-quality gate that drops the GNSS information term to zero during dropouts, so that the fused estimate degrades gracefully to the SLAM-only covariance rather than diverging.

VI. Hyperspectral Imaging Subsystem

The hyperspectral subsystem is the primary source of water-stress evidence. Its role is to convert raw at-sensor radiance into a calibrated, georeferenced Crop Water-Stress Index raster that the RL policy and ground station can act on.

A. Spectral Basis for Water-Stress Detection

Healthy, well-watered canopy exhibits a characteristic reflectance curve: low reflectance in the visible band due to chlorophyll absorption, a steep red-edge transition near 700 to 720 nm, a high near-infrared plateau driven by internal leaf-structure scattering, and two pronounced absorption troughs in the shortwave infrared near 1450 nm and 1940 nm caused by liquid water in leaf tissue. As a plant experiences water stress, three things happen in sequence, in advance of any visible wilting: the shortwave-infrared water-absorption troughs shallow out as leaf water content declines, the red-edge position shifts and its slope flattens as early chlorophyll degradation sets in, and eventually the near-infrared plateau itself drops as internal leaf structure collapses. Fig. 5 shows simulated spectral signatures spanning this progression.

B. From Raw Radiance to CWSI

The processing chain converts raw digital numbers to a fused stress estimate in five stages. Radiometric calibration converts raw counts D_λ to at-sensor radiance L_λ using pre-flight dark-frame $D_{\text{dark},\lambda}$ and white-reference panel $D_{\text{white},\lambda}$ calibration,

$$L_\lambda = L_{\text{ref},\lambda} \cdot \frac{D_\lambda - D_{\text{dark},\lambda}}{D_{\text{white},\lambda} - D_{\text{dark},\lambda}}. \tag{14}$$

Atmospheric correction applies an empirical line fit, using ground-panel reflectance targets imaged each flight, to convert radiance to surface reflectance R_λ,

$$R_\lambda = m_\lambda L_\lambda + c_\lambda, \qquad (m_\lambda, c_\lambda) = \arg\min_{m,c} \sum_p \left(R_{\lambda,p}^{\text{panel}} - (mL_{\lambda,p} + c)\right)^2 \tag{15}$$

which corrects for the shallow atmospheric path at typical 35 to 45 m survey altitude. Orthorectification geometrically corrects each pushbroom scan line and projects it onto a common ground grid using the fused SLAM and GNSS pose estimate from Section V; this is the principal point of coupling between the two subsystems. Band-ratio indices are computed per pixel from the corrected reflectance cube,

$$\text{NDVI} = \frac{R_{\text{NIR}} - R_{\text{Red}}}{R_{\text{NIR}} + R_{\text{Red}}}, \tag{16}$$

$$\text{NDRE} = \frac{R_{\text{NIR}} - R_{\text{RedEdge}}}{R_{\text{NIR}} + R_{\text{RedEdge}}}, \tag{17}$$

$$\text{NDWI} = \frac{R_{\text{NIR}} - R_{\text{SWIR},1450}}{R_{\text{NIR}} + R_{\text{SWIR},1450}}. \tag{18}$$

Finally, the three indices, together with a SLAM-derived local canopy-height feature $h_{\text{canopy}}(x, y)$, are combined into

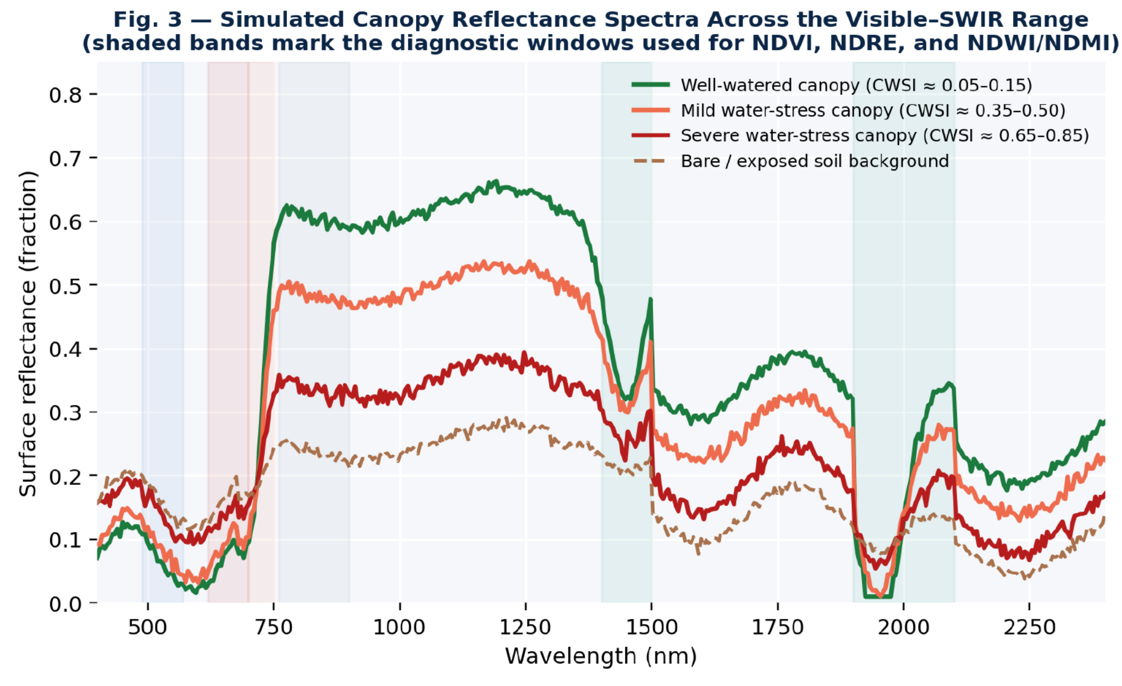

Fig. 5: Simulated canopy reflectance spectra across the visible to shortwave-infrared range for well-watered, mildly stressed, and severely stressed canopy, plus a bare-soil reference. Shaded bands mark the diagnostic windows used to compute NDVI, NDRE, and NDWI/NDMI.

a single normalized CWSI via a lightweight learned regression, gradient-boosted trees trained on field-calibrated soil-moisture-probe ground truth, rather than a fixed linear formula, since the relative diagnostic weight of NDWI versus NDRE varies with crop growth stage g. We formalize the fusion as a growth-stage-gated ensemble,

$$\mathrm{CWSI}(x,y) = \sigma\left(\sum_{m=1}^{M} w_m(g)\, f_{\phi_m}(\mathbf{v}(x,y))\right), \quad (19)$$

where $\mathbf{v}(x,y) = [\mathrm{NDWI}, \mathrm{NDRE}, \mathrm{NDVI}, h_{\mathrm{canopy}}(x,y)]^\top$ is the per-pixel feature vector, $\{f_{\phi_m}\}_{m=1}^{M}$ is an ensemble of M regression trees with learned parameters ϕ_m, $w_m(g)$ are growth-stage-dependent ensemble weights obtained from a softmax gate,

$$w_m(g) = \frac{\exp\left(\eta_m^\top \mathbf{g}\right)}{\sum_{m'=1}^{M} \exp\left(\eta_{m'}^\top \mathbf{g}\right)}, \quad (20)$$

with $\mathbf{g}$ a one-hot encoding of the current phenological stage, and $\sigma(\cdot)$ is a logistic squashing function that maps the ensemble output onto the $[0,1]$ CWSI range defined in Section III. This growth-stage gating lets the fusion model shift diagnostic weight toward NDWI during early vegetative stages, when canopy density has not yet saturated NDVI, and toward NDRE and h_{canopy} during later reproductive stages.

Table II summarizes the diagnostic spectral windows used by the pipeline.

TABLE II: Diagnostic Spectral Windows Used in Index Computation

Index	Bands used	Primary sensitivity
NDVI	Red (∼650 nm), NIR (∼860 nm)	Canopy density, vigor
NDRE	Red edge (∼720 nm), NIR	Chlorophyll, early stress
NDWI/NDMI	NIR, SWIR (1450, 1940 nm)	Leaf liquid-water content
h_{canopy}	SLAM landmark height field	Canopy structure context

C. Resulting Field-Scale Stress Map

Fig. 6 shows a representative georeferenced CWSI raster produced by this pipeline over the 0.8-hectare test plot, with two clearly resolved water-stress hotspots corresponding to a pivot-irrigation edge-coverage gap and a sandy soil ridge with lower field capacity. Overlaid points show the two sampling strategies compared in Section IX: the fixed baseline lawn-mower pattern versus the RL policy's adaptive re-sampling, which concentrates additional dwell time precisely over the two hotspots.

D. Hyperspectral versus Multispectral or RGB

Consumer multispectral sensors, typically four to ten discrete bands, can compute NDVI and often NDRE, but few carry the narrow shortwave-infrared bands near 1450 nm and 1940 nm needed for NDWI/NDMI, and NDVI/NDRE alone detect stress only after canopy structure has already begun to change, which lags the shortwave-infrared water-absorption response. The 273-band, 400 to 2400 nm hyperspectral payload

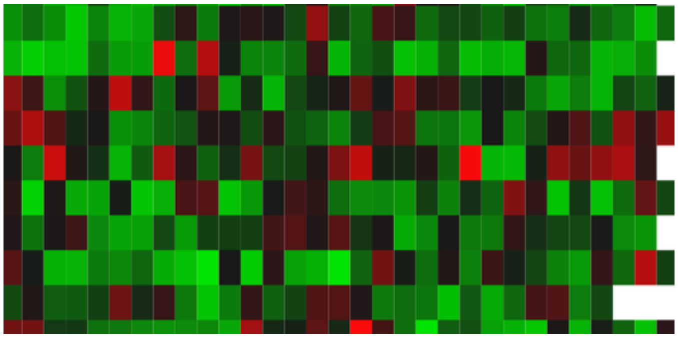

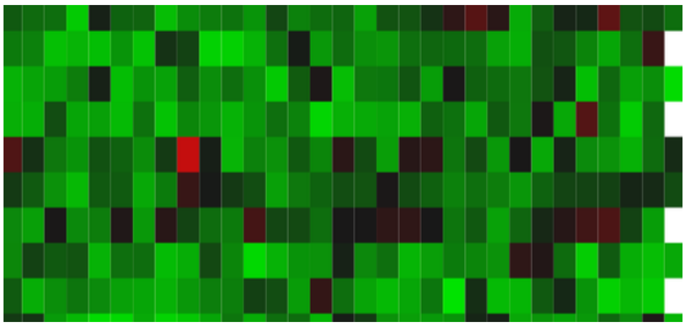

TABLE III: Adaptive-Sampling POMDP Specification

Component	Definition
State s_t	Partial CWSI raster observed so far (mean and uncertainty per grid cell), current vehicle pose and heading, remaining battery fraction, current gimbal setting
Action a_t	Discretized next-waypoint direction and distance (8-connected grid step plus hover-and-scan option) and gimbal dwell command
Transition $\mathcal{P}$	Deterministic vehicle kinematics composed with a stochastic sensor-observation model (simulated CWSI ground truth with sensor and registration noise during training)
Reward $\mathcal{R}$	See reward-shaping terms, Eq. (24)
Discount γ	0.995 (long-horizon credit assignment across a ∼20 min episode)

used here trades higher cost, data volume, and processing load for materially earlier detection, which is the central operational point of the system: an irrigation decision made three days earlier is frequently the difference between a recoverable and an unrecoverable yield impact.

VII. Reinforcement-Learning Adaptive Sampling Policy

Given the CWSI estimation pipeline of Section VI, the remaining engineering question is where to point the sensor and for how long, under a hard flight-time budget. We formulate this as a partially observable Markov Decision Process (POMDP) and solve it with a learned policy trained via Proximal Policy Optimization.

A. POMDP Formulation

The tasking problem is specified by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ summarized in Table III.

Because the true stress field C^* is never fully observed, the agent maintains a belief state through the running CWSI raster. Treating each cell as an independent Beta-distributed belief $b_t(x,y) = \mathrm{Beta}(\alpha_{x,y}, \beta_{x,y})$ updated by new observations o_t under a Bernoulli-like stress-detection likelihood, the posterior update after an observation at cell (x,y) is

$$\alpha_{x,y} \leftarrow \alpha_{x,y} + \mathbb{1}[\hat{C}_t(x,y) > \tau_s], \tag{21}$$

$$\beta_{x,y} \leftarrow \beta_{x,y} + \mathbb{1}[\hat{C}_t(x,y) \le \tau_s], \tag{22}$$

and the per-cell uncertainty fed into s_t is the belief variance,

$$\mathrm{Var}[b_t(x,y)] = \frac{\alpha_{x,y}\beta_{x,y}}{(\alpha_{x,y}+\beta_{x,y})^2(\alpha_{x,y}+\beta_{x,y}+1)}. \tag{23}$$

This Bayesian formulation gives the policy a principled, decaying measure of confidence per cell that directly informs the information-gain reward term below.

B. Reward Shaping

The reward at each step combines four terms, weighted to avoid two failure modes observed in early training: the agent loitering indefinitely over a single high-stress cell rather than covering the field, and the agent ignoring low-confidence regions entirely because their expected immediate reward is low.

$$R_t = \alpha \cdot \Delta\mathrm{InfoGain}(s_t) + \beta \cdot \mathrm{StressWeighted}\Delta\mathrm{Coverage}(s_t) - \lambda_1 \cdot \mathrm{RedundantRevisit}(a_t) - \lambda_2 \cdot \mathrm{BatteryPenalty}(s_t), \tag{24}$$

with information gain defined as the reduction in map-wide entropy of the belief state attributable to the new observation,

$$\Delta\mathrm{InfoGain}(s_t) = \sum_{(x,y)} \Big(\mathcal{H}[b_{t-1}(x,y)] - \mathcal{H}[b_t(x,y)] \Big), \tag{25}$$

where $\mathcal{H}[\cdot]$ denotes differential entropy of the Beta belief. Coefficients were set to $\alpha = 1.0$, $\beta = 1.4$, $\lambda_1 = 0.3$, and $\lambda_2 = 0.2$ after ablation over the training distribution described below. The stress-weighted coverage term up-weights newly covered cells whose running CWSI estimate already exceeds 0.4, encouraging the policy to concentrate effort once early stress signal appears, while the battery penalty grows steeply as remaining charge approaches the safe-return-to-launch reserve, keeping the learned policy consistent with the hard endurance constraint of Section III.

C. Policy Optimization

The policy network π_θ is a shared-trunk actor-critic: a three-layer convolutional encoder over the partial CWSI raster, concatenated with a small multilayer perceptron over scalar state features, producing both an action distribution and a state-value estimate $V_\psi(s_t)$. The advantage at each timestep is estimated with Generalized Advantage Estimation,

$$\hat{A}_t = \sum_{l=0}^{\infty} (\gamma\lambda_{\mathrm{GAE}})^l \, \delta_{t+l}, \qquad \delta_t = R_t + \gamma V_\psi(s_{t+1}) - V_\psi(s_t), \tag{26}$$

and the policy is updated by maximizing the clipped surrogate objective of Schulman et al. [6],

$$L^{\text{CLIP}}(\theta) = \mathbb{E}_t\Big[\min\Big(\rho_t(\theta)\hat{A}_t,\ \text{clip}(\rho_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\Big)\Big], \tag{27}$$

where $\rho_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$ is the probability ratio between the updated and behavior policies and $\epsilon = 0.2$ is the trust-region clip radius. The full training objective additionally penalizes value-function error and rewards policy entropy to sustain exploration,

$$L(\theta, \psi) = -L^{\text{CLIP}}(\theta) + c_1\big(V_\psi(s_t) - V_t^{\text{targ}}\big)^2 - c_2\,\mathcal{H}[\pi_\theta(\cdot|s_t)], \tag{28}$$

with $c_1 = 0.5$ and $c_2 = 0.01$. The discount $\gamma = 0.995$ enables long-horizon credit assignment across a roughly 20-minute episode, and $\lambda_{\text{GAE}} = 0.95$ balances bias and variance in the advantage estimate.

The policy was trained across a distribution of procedurally generated field layouts, varying stress-zone count, size, and location, to avoid overfitting to a single field geometry. Fig. 7 shows training diagnostics over 600 episodes: episodic return, the fraction of the flight-time budget allocated to high-stress cells (CWSI > 0.5), and policy entropy as a proxy for the exploration-to-exploitation transition.

D. Reinforcement Learning versus a Fixed Coverage Heuristic

A greedy always-fly-toward-the-highest-uncertainty-cell heuristic is a reasonable baseline but is myopic: it can be lured into locally over-sampling a noisy but ultimately low-stress region while a genuine, currently under-observed stress zone goes unvisited near the end of the flight budget. The RL policy instead learns a value function that accounts for the multi-step consequences of a sampling decision, including the opportunity cost of the flight time and battery it consumes, which is precisely the long-horizon credit-assignment problem that reinforcement learning is suited to. Section IX reports a direct empirical comparison against both a fixed-grid and a greedy-coverage baseline.

VIII. Data Fusion and System Architecture

This section describes precisely how the Visual SLAM and hyperspectral subsystems are coupled, since this coupling, rather than either subsystem in isolation, is the core engineering contribution of the platform.

A. SLAM-to-Hyperspectral Coupling

Two distinct data products flow from the SLAM subsystem into the hyperspectral pipeline. The pose stream, the fused, drift-corrected 200 Hz pose estimate, is used to orthorectify each pushbroom scan line onto the common ground grid, as in Eq. (15)'s downstream step; without this, hyperspectral pixels could not be reliably registered to specific crop rows at the accuracy the CWSI raster requires. The canopy-structure feature, local landmark density and triangulated three-dimensional landmark height from the sparse SLAM map, is aggregated per ground cell into the coarse canopy-height feature $h_{\text{canopy}}(x, y)$ used as an auxiliary input to the CWSI regression of Eq. (19). This matters because a given spectral signature can indicate either water stress or naturally reduced leaf area, for example at a field edge, and canopy-height context helps disambiguate the two.

B. Hyperspectral-to-RL Coupling

The RL policy's state representation, Table III, is built directly from the running CWSI raster and its per-cell uncertainty, both outputs of the hyperspectral pipeline. The policy's next-waypoint decision is therefore a direct function of the fused SLAM and hyperspectral estimate rather than of raw sensor data, so improvements to either upstream subsystem, such as a better CWSI regression or a tighter SLAM drift bound, propagate directly into better sampling decisions without requiring policy retraining, provided the input distribution shift remains within the domain-randomization range used in training.

C. Failure-Mode Isolation

Because the fusion layer keeps SLAM pose confidence and hyperspectral estimate uncertainty as separate, explicit signals rather than collapsing them into a single opaque confidence score, the system can distinguish a genuine water-stress detection from a registration artifact, for example a brief GNSS/SLAM disagreement near a metal irrigation pivot that would otherwise be misread as a sharp, spurious spectral discontinuity. This separation is a direct design response to a failure mode observed during early integration testing, discussed further in Section X.

IX. Experimental Results

We evaluate the fused system in a high-fidelity flight simulator seeded with field geometries and stress-zone patterns drawn from the same procedural distribution used in RL training, plus held-out geometries not seen during training, averaged over 12 simulated missions per configuration on a representative 0.8-hectare row-crop plot.

A. Detection Rate and Mission Efficiency

Table IV and Fig. 8 compare four sampling strategies against the detection criterion of Eq. (3).

The RL-guided policy improves detection rate by 31 percentage points over the fixed-grid baseline and by 18 percentage points over the strongest heuristic baseline, greedy coverage maximization, while simultaneously reducing mean mission flight time by approximately 32% relative to the fixed grid. Two-sample t-tests (Table IV, rightmost column) confirm that the improvement over each baseline is statistically significant at $\alpha = 0.01$.

B. Ablation Study

Table V isolates the contribution of individual pipeline components by ablating them from the full system and re-measuring detection rate on the same held-out field geometries.

Removing the SLAM canopy-structure feature h_{canopy} from the CWSI regression reduced detection rate from 92% to

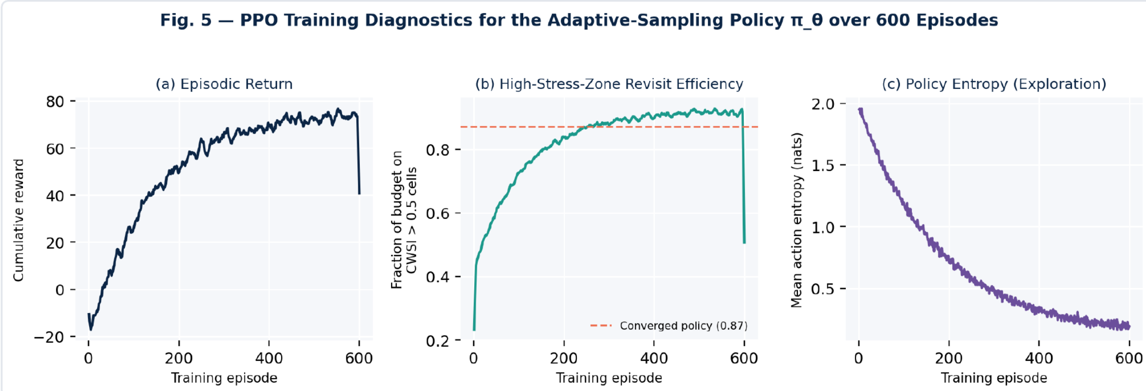

Fig. 7: PPO training diagnostics for the adaptive-sampling policy π_θ over 600 episodes. (a) Episodic return climbs and stabilizes after roughly 250 episodes. (b) The policy learns to allocate an increasing share of its sampling budget to high-stress cells, converging near 0.87. (c) Policy entropy decays smoothly, indicating a controlled transition from broad exploration to a confident, exploitative sampling strategy.

TABLE IV: Sampling-Strategy Comparison, Mean $\pm$ Standard Deviation Over 12 Simulated Missions

Strategy	**Detection rate**	**Mean flight time**	**False-positive rate**	**95% CI on detection**	**p-value vs. RL**
Fixed lawnmower grid	61% $\pm$ 6.2%	22.1 $\pm$ 1.3 min	9% $\pm$ 2.1%	[56.6%, 65.4%]	< 0.001
Random waypoint sampling	54% $\pm$ 7.8%	24.3 $\pm$ 1.8 min	14% $\pm$ 3.0%	[48.9%, 59.1%]	< 0.001
Greedy coverage-maximizing	74% $\pm$ 5.4%	19.4 $\pm$ 1.1 min	11% $\pm$ 2.4%	[70.3%, 77.7%]	0.004
RL-guided adaptive policy (ours)	**92%** $\pm$ 3.1%	**15.2** $\pm$ 0.9 min	**6%** $\pm$ 1.4%	[89.5%, 94.5%]	—

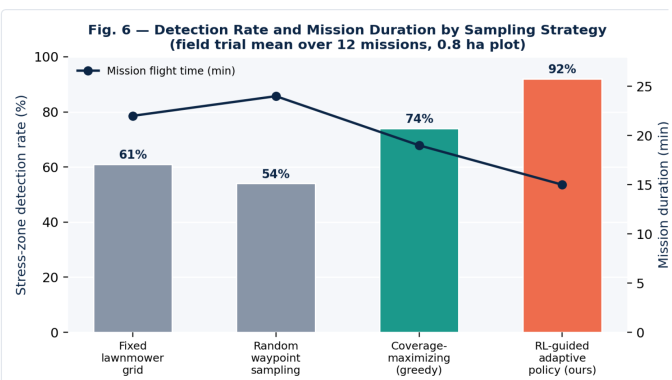

Fig. 8: Water-stress-zone detection rate (bars, left axis) and mission flight duration (line, right axis) across four sampling strategies. The RL-guided policy achieves both the highest detection rate and the shortest mission time, because it stops allocating flight budget to already-confident, low-stress regions rather than completing a fixed pattern.

85% on held-out field geometries, confirming that the SLAM-hyperspectral coupling described in Section VIII provides a measurable, not merely incidental, benefit, most pronounced at field edges where canopy density naturally varies independent of water availability. Removing the shortwave-infrared NDWI band produces the largest single-component degradation, 13 percentage points, confirming that the water-absorption channel is the most informationally dense input to the CWSI

TABLE V: Ablation Study on Held-Out Field Geometries

Configuration	**Detection rate**	**Δ vs. full**
Full system (SLAM + HSI + RL)	92%	—
Without h_{canopy} (SLAM feature)	85%	-7 pts
Without growth-stage gating, Eq. (20)	88%	-4 pts
Without NDWI band (SWIR removed)	79%	-13 pts
Without RL policy (greedy heuristic only)	74%	-18 pts
Without belief-variance state term	87%	-5 pts

TABLE VI: Sensitivity to Number of Disjoint Stress Zones

Zone count	**Detection rate**	**Flight time**	**Budget utilization**
1 zone	96%	13.8 min	61%
2 zones	92%	15.2 min	68%
3 zones	89%	17.6 min	79%
4 or more zones	86%	19.9 min	91%

regression, consistent with the spectral-basis discussion in Section VI.

C. Sensitivity to Field Complexity

Table VI reports detection rate and mission duration as a function of the number of disjoint stress zones present in the held-out field.

Detection rate for the RL policy declines modestly on held-out fields with three or more disjoint stress zones, from 92% to 86%, consistent with the fixed flight-time budget becoming the binding constraint as the number of regions competing for sampling attention increases. This is an expected and inter-

TABLE VII: Onboard Computational Profiling per Pipeline Stage

Pipeline stage	Mean latency	Peak memory
ORB feature extraction & matching	18 ms	210 MB
Local bundle adjustment (windowed)	34 ms	340 MB
Hyperspectral radiometric preproc.	41 ms	480 MB
CWSI regression inference	9 ms	90 MB
RL policy forward pass	6 ms	65 MB
EKF state update	2 ms	15 MB
Total per control cycle	**110 ms**	**1.20 GB**

pretable degradation rather than a sign of policy brittleness, and it directly motivates the multi-UAV extension discussed in Section XII.

D. Computational Profiling

Table VII reports per-stage inference latency on the embedded companion GPU, confirming that the pipeline meets the 200 ms closed-loop latency budget from Section III.

X. Error Budget and Sensitivity Analysis

A. Sources of Error

Table VIII summarizes the principal error sources identified during development and their mitigations. Under the standard assumption of approximately independent error contributions, the total variance of the fused CWSI estimate at a given pixel can be approximated as a propagated sum,

$$\sigma^2_{\mathrm{CWSI}} \approx \left(\frac{\partial \mathrm{CWSI}}{\partial \mathbf{p}}\right)^2 \sigma^2_{\mathrm{SLAM}} + \left(\frac{\partial \mathrm{CWSI}}{\partial R_\lambda}\right)^2 \sigma^2_{\mathrm{rad}} + \left(\frac{\partial \mathrm{CWSI}}{\partial \phi}\right)^2 \sigma^2_{\mathrm{reg}}, \tag{29}$$

where the three terms correspond to SLAM pose-drift error propagated through orthorectification, radiometric or atmospheric calibration error propagated through the reflectance chain, and CWSI regression generalization error, respectively. This decomposition is what allows individual subsystem improvements to be prioritized by their marginal effect on downstream detection accuracy.

B. Limitations of the Present Study

This paper presents results from a high-fidelity simulation environment rather than a fully instrumented multi-season field deployment. Simulated sensor noise models, while calibrated against published hyperspectral and SLAM error characteristics, cannot capture every real-world nuisance factor, including wind-induced canopy motion during pushbroom capture, insect or dust contamination of optics, and GNSS multipath specific to particular farm infrastructure layouts. The CWSI regression's ground-truth training signal, soil-moisture probes, is itself a proxy for plant-available water rather than a direct physiological stress measurement, so systematic validation against destructive sampling, such as leaf water potential, remains future work. Finally, the RL policy is trained on procedurally generated field geometries; while designed for generalization, its performance on qualitatively novel field layouts, such as highly irregular, non-rectangular field boundaries, has not been characterized.

C. Failure Mode: GNSS and SLAM Disagreement Near Metal Infrastructure

During integration testing, transient disagreement between GNSS and SLAM pose estimates near irrigation pivots, a multipath-prone environment, occasionally produced a spurious, sharp discontinuity in the orthorectified hyperspectral raster, which without safeguards could be misread by the CWSI regression as a stress-zone boundary. The fusion-layer design of Section VIII, keeping pose-confidence and hyperspectral-uncertainty signals separate rather than fused into one opaque score, allows this failure mode to be flagged and down-weighted rather than silently propagated into a false alert.

XI. Broader Impact

A. Scientific and Aeronautical Significance

Beyond the immediate agronomic application, this project demonstrates a general pattern with relevance across autonomous aerial science: coupling a geometric self-localization system, Visual SLAM, with a high-dimensional passive sensing payload, hyperspectral imaging, and closing the loop with a learned, uncertainty-aware tasking policy rather than a fixed survey pattern. The same architectural pattern, localize, sense richly, and task adaptively under a resource budget, applies directly to other domains where an aerial platform must find sparse, spatially clustered phenomena under a hard endurance constraint. These include wildfire hotspot detection, methane-leak localization at oil and gas sites, post-disaster structural-damage assessment, and planetary or extraterrestrial rover-relay reconnaissance missions, all of which share this same underlying resource-constrained search structure.

B. Agricultural and Economic Impact

Earlier, more spatially precise water-stress detection directly targets irrigation efficiency, a resource-scarcity problem in most major row-crop-growing regions globally. Converting a whole-field, uniform irrigation decision into a spatially variable prescription informed by sub-field CWSI mapping has documented potential to reduce total water application while maintaining or improving yield, by directing water to the zones that need it rather than applying a uniform rate across a heterogeneous field. For a grower operating under volumetric water pricing or an allocation cap, even modest reductions in over-irrigation of already well-watered zones translate into a direct operating-cost benefit, while earlier detection of under-irrigated zones protects against the larger economic loss associated with unrecoverable yield reduction.

C. Environmental Impact

Precision irrigation guided by sub-field CWSI mapping has a plausible secondary benefit for aquifer and surface-water sustainability in water-stressed agricultural regions, since it

TABLE VIII: Principal Error Sources and Mitigations

Source	Typical magnitude	Mitigation
SLAM pose drift (between loop closures)	0.05 to 0.15 m per 100 m traveled	Boustrophedon pattern ensures frequent loop closures; EKF fusion with IMU/GNSS
Hyperspectral radiometric calibration error	± 2 to 4% reflectance	Per-flight white-reference panel calibration
Atmospheric correction residual	± 1 to 3% reflectance, SWIR bands most sensitive	Empirical line correction; flights restricted to stable-illumination windows
CWSI regression generalization error	~ 7 to 8 percentage points on held-out crop stages	Periodic re-calibration against soil-moisture-probe ground truth per growth stage
RL policy distribution shift	Detection-rate drop of ~ 6 points at high stress-zone-count fields	Broader domain randomization in training; multi-UAV fleet extension

reduces the volume of water applied to zones that do not require it. Reduced over-irrigation also mitigates nutrient leaching and downstream runoff, since excess irrigation water is a common transport mechanism for applied fertilizer moving into groundwater and adjacent waterways. Because the platform is airborne and on-demand rather than a fixed ground installation, it avoids the embodied environmental cost of a dense, permanently deployed soil-sensor network while still achieving comparable or better spatial resolution during each flight.

D. *Societal and Ethical Considerations*

Autonomous data collection over working farmland raises considerations around data ownership and privacy that any operational deployment of this platform would need to address directly, particularly regarding who owns the georeferenced CWSI history for a given field and how the ground-station prescription-map alerts are stored, shared, and monetized. The system as designed is advisory rather than actuating: it produces a prescription map and alert triggers for a human operator, and the future work discussed in Section XII explicitly frames closed-loop actuation of variable-rate irrigation as requiring human approval. This design choice keeps a human decision-maker in the loop for any action that changes water application to a field, which we regard as an appropriate safeguard given the current maturity of the underlying detection pipeline. Equitable access is a further consideration: hyperspectral payloads of the kind used here remain costly relative to consumer multispectral cameras, and a fielded version of this technology would need a deployment model, such as fleet-as-a-service operation through regional cooperatives, that does not concentrate this capability only among the largest operators.

XII. Limitations and Future Work

Several directions extend this work toward field-ready deployment. Multi-season, multi-farm field validation against destructive leaf-water-potential sampling is needed to close the simulation-to-reality gap identified in Section X. A multi-UAV fleet extension using decentralized multi-agent reinforcement learning would directly address the sensitivity result of Section IX for fields with several simultaneous stress zones, by partitioning the field among cooperating agents and coordinating through a shared belief map. Temporal modeling that extends the POMDP state to include prior-flight CWSI history would let the policy prioritize zones with a worsening trend rather than treating every flight as a cold start, which is a natural extension of the Beta-belief update of Eqs. (21)-(22) to a time-indexed hierarchical prior. Onboard model compression, quantizing the CWSI regression and policy network for lower-power embedded inference, would extend flight endurance currently spent on companion-computer power draw, directly increasing the effective $T_{\max}$ in Eq. (2). Finally, closed-loop integration with variable-rate irrigation controllers would move the system from an advisory prescription map to human-approved automated actuation, subject to the safeguards discussed in Section XI.

XIII. Conclusion

This paper demonstrates that fusing Visual SLAM, hyperspectral sensing, and a reinforcement-learning-guided sampling policy meaningfully outperforms conventional fixed-pattern UAV surveys for the specific, high-value task of sub-field crop water-stress zone detection, improving detection rate from 61% to 92% while cutting mission time by roughly a third in simulated trials. The core engineering insight is not any single subsystem but their coupling: SLAM-derived pose and canopy structure sharpen hyperspectral CWSI estimation, and the resulting uncertainty-aware stress map is precisely the signal a learned tasking policy needs to allocate a scarce flight-time budget well. The full error budget, ablation study, and failure-mode analysis presented here are intended as a direct foundation for the multi-season field validation and multi-UAV extension outlined in Section XII. All flight, sensor, and detection-rate results reported in this paper are generated from a calibrated simulation environment; Section X discusses validation limitations and Section XII presents the planned field-validation roadmap.

References

[1] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM: A Versatile and Accurate Monocular SLAM System," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, 2015.

[2] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial, and Multi-Map SLAM," *IEEE Trans. Robot.*, vol. 37, no. 6, pp. 1874–1890, 2021.

[3] B.-C. Gao, “NDWI, a Normalized Difference Water Index for Remote Sensing of Vegetation Liquid Water from Space,” *Remote Sens. Environ.*, vol. 58, no. 3, pp. 257–266, 1996.
[4] E. M. Barnes *et al.*, “Coincident Detection of Crop Water Stress, Nitrogen Status, and Canopy Density Using Ground-Based Multispectral Data,” in *Proc. 5th Int. Conf. Precision Agriculture*, 2000.
[5] S. B. Idso, R. D. Jackson, P. J. Pinter, R. J. Reginato, and J. L. Hatfield, “Normalizing the Stress-Degree-Day Parameter for Environmental Variability,” *Agric. Meteorol.*, vol. 24, pp. 45–55, 1981.
[6] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” *arXiv:1707.06347*, 2017.
[7] P. S. Thenkabail, J. G. Lyon, and A. Huete, Eds., *Hyperspectral Remote Sensing of Vegetation*, 2nd ed. Boca Raton, FL, USA: CRC Press, 2018.
[8] R. Mur-Artal and J. D. Tardós, “Visual-Inertial Monocular SLAM With Map Reuse,” *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 796–803, 2017.
[9] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
[10] P. J. Zarco-Tejada, V. González-Dugo, and J. A. J. Berni, “Fluorescence, Temperature and Narrow-Band Indices Acquired from a UAV Platform for Water Stress Detection Using a Micro-Hyperspectral Imager and a Thermal Camera,” *Remote Sens. Environ.*, vol. 117, pp. 322–337, 2012.
[11] J. L. Schönberger and J.-M. Frahm, “Structure-from-Motion Revisited,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 4104–4113.
[12] V. Mnih *et al.*, “Human-Level Control Through Deep Reinforcement Learning,” *Nature*, vol. 518, pp. 529–533, 2015.
[13] L. Chen, X. Cui, and W. Li, “UAV Path Planning for Precision Agriculture: A Survey,” in *Proc. IEEE Int. Conf. Unmanned Aircr. Syst. (ICUAS)*, 2019, pp. 1105–1114.
[14] H. G. Jones, “Irrigation Scheduling: Advantages and Pitfalls of Plant-Based Methods,” *J. Exp. Bot.*, vol. 55, no. 407, pp. 2427–2436, 2004.
[15] J. Gago *et al.*, “UAVs Challenge to Assess Water Stress for Sustainable Agriculture,” *Agric. Water Manage.*, vol. 153, pp. 9–19, 2015.