Overcome the Fear Of Missing Out: Active Sensing UAV Scanning for Precision Agriculture

Marios Krestenitis[a], Emmanuel K. Raptis[b,a], Athanasios Ch. Kapoutsis[a], Konstantinos Ioannidis[a], Elias B. Kosmatopoulos[b,a], Stefanos Vrochidis[a]

[a] *Information Technologies Institute, Centre for Research & Technology Hellas, Thessaloniki, 57001, Greece*
[b] *Department of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, 67100, Greece*

Abstract

This paper deals with the problem of informative path planning for a UAV deployed for precision agriculture applications. First, we observe that the "fear of missing out" data lead to uniform, conservative scanning policies over the whole agricultural field. Consequently, employing a non-uniform scanning approach can mitigate the expenditure of time in areas with minimal or negligible real value, while ensuring heightened precision in information-dense regions. Turning to the available informative path planning methodologies, we discern that certain methods entail intensive computational requirements, while others necessitate training on an ideal world simulator. To address the aforementioned issues, we propose an active sensing coverage path planning approach, named OverFOMO, that regulates the speed of the UAV in accordance with both the relative quantity of the identified classes, i.e. crops and weeds, and the confidence level of such detections. To identify these instances, a robust Deep Learning segmentation model is deployed. The computational needs of the proposed algorithm are independent of the size of the agricultural field, rendering its applicability on modern UAVs quite straightforward. The proposed algorithm was evaluated with a simu-realistic pipeline, combining data from real UAV missions and the high-fidelity dynamics of AirSim simulator, showcasing its performance improvements over the established state of affairs for this type of missions. An open-source implementation of the algorithm and the evaluation pipeline is also available: `https://github.com/emmarapt/OverFOMO`.

Keywords: Adaptive Path Planning, Semantic Segmentation, UAV Imagery, Precision Agriculture, AirSim

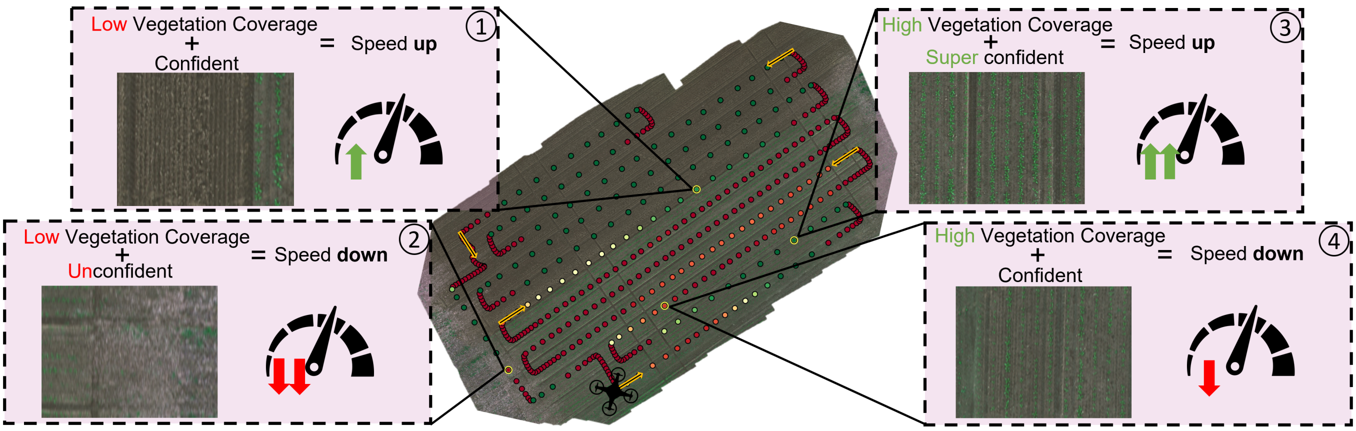

Figure 1: Graphical illustration of the core rationale behind the proposed active sensing approach. UAV is scanning a field with an on board system to estimate the vegetation coverage via captured images, the objective is to on-line regulate its speed so as 1) to cover in detail the whole area and 2) in the minimum possible time. Intuitively, one would like to speed up in areas with little to no information, i.e. vegetation coverage, (Snapshot 1) and slowdown in areas that have rich information to be sure that it can capture everything in great detail (Snapshot 4). However, the amount of information is not the sole factor that should define such changes, as the system that estimates this information could be occasionally inaccurate, mostly due to camera movement. In Snapshot 2, although probably there is not significant information underneath, the UAV should slow down to increase its confidence and be sure about this estimation. On the other hand, Snapshot 3 illustrates a case where, although the vegetation coverage is definitely high, the absolute certainty in such estimation allows for an extra increase in the UAV speed, allowing to save precious flight time.

1. Introduction

Unmanned Aerial Vehicles (UAVs) are probably the robotics platforms with the highest adoption rate from professionals in their fields. For example, UAVs are now vital assets for rescuers to quickly search large areas [1], construction engineers to monitor their project's evolution [2], firefighters to quickly assess and identify the fire front [3], farmers and agronomist to effectively assess the crops health [4], etc. Such a diverse adoption drives more people and effort to be devoted to UAVs related research and development, leading, in turn, to a further increase in the type of the supported applications. One of the critical factors that have affected this UAV success cycle is the recent advancements in deep learning and specifically in computer vision tasks [5]. Now, more than ever, we have at our disposal powerful tools that can process the UAV-related data both in an offline and onboard fashion. One of the most severe bottlenecks has to do with the available quantity and quality of data (diverse, clean, and annotated) to deploy the deep learning techniques. Therefore, it is of paramount importance to develop efficient methodologies for automatic meaningful data acquisition, using limited infrastructures.

One of the UAV application areas that could benefit greatly from such methodologies lies within the precision agriculture domain. More precisely, in these applications the UAV collected data are used, in post-processing fashion, to construct homogeneous orthomosaic [6], define the crops' health [7], find crop line [8], detect and recognize harmful weeds [9], etc. The problem to be investigated in this paper deals with the intelligent design of UAV scanning policy, so as to avoid spending time in areas with little to no real value while being extra precise in information-rich areas. In essence, we seek to answer the following question: Can the online received information "steer" the UAV towards a more efficient data collection policy? In literature, this problem is usually referred to as Informative Path Planning (IPP) [10, 11].

1.1. Related Work

Currently, the vast majority of the UAV agriculture coverage mission planners applies a variance of back-and-forth methodology exploiting the Spanning-Tree Coverage (STC) algorithm [12], or boustrophedon approach [13]. Although this family of approaches is relatively simple, it has been proven quite effective, rendering it the "go-to" approach [14]. The problem with such approaches is the implied assumption of a uniform distribution of the information across the field to be surveyed. In practice, this is rarely the case, forcing the UAV path to be either *too pessimistic* and eventually cover fewer square meters than it could or *not pessimistic enough* resulting in inadequate coverage in specific subparts of the field. Recognizing that, a fair amount of IPP works have been proposed ,which deploy a trajectory adjusting mechanism based on the online received data.

Research-wise, a large number of UAV-based IPP applications have been developed using Gaussian processes (GPs) as a natural way of encoding spatial correlations among the online received data and creating terrain maps of continuous scalar fields. Within the realm of IPP, GPs have gained considerable popularity as a Bayesian method for effectively modeling spatiotemporal phenomena and their inherent correlations [15], enabling the collection of data that takes into account both map structure and uncertainty. However, the primary challenge encountered when directly applying Gaussian Processes (GPs) [16, 17] is the significant computational burden that arises due to the accumulation of dense imagery data over time.

Ruckin et al. [10] introduced an IPP methodology utilizing Bayesian techniques as an active learning acquisition function to quantify the pixel-wise model uncertainty in semantic segmentation. Their approach aimed to maximize the improvement in the model's performance by assimilating the most informative terrain data with the highest uncertainty, linking thus the information gain from the active learning acquisition function to a planning objective. Vivaldini et al. [17] proposed an online UAV-based IPP system, wherein the acquisition function is designed to minimize the uncertainty associated with the differentiation between diseased trees and healthy trees as well as roads in a Gaussian map interpolation. The path planning module strategically selects sampling points to achieve comprehensive environmental coverage, utilizing the Rapidly-exploring Random Trees (RTT) algorithm to optimize the gathering of crucial information. To minimize the distance traveled and ensure sufficient coverage of the surveyed area, an objective function is responsible for guiding the UAV toward reducing the average uncertainty of an image at a given position (x, y) on the current classification map. Although their experimental evaluation demonstrated favorable outcomes in comparison to static coverage paths, the proposed methodology allocates the UAV's battery life to repetitive back-and-forth movements, which undeniably leads to suboptimal efficiency in continuous terrain monitoring.

Popovic et al. [18, 19] proposed an IPP framework for active classification, exploiting the spatial correlation encoded in a Gaussian Process model as a prior for Bayesian data fusion to facilitate expedited map updates. They proposed an adaptable path-planning approach that generates dynamically viable trajectories at varying altitudes in a continuous 3D space to achieve high-quality aerial imaging with constant-time measurements by computing the informative objective with the new map representation. Their strategy, however, assumes swift map updates with minimal computational overhead, while simultaneously allocating the UAV's temporal resources to vertical maneuvers. While their simulated and real-life experiments yielded positive results when compared to static coverage paths, it is noteworthy that the suggested methodology has predominantly been appraised in limited-scale field trials where the temporal exigency of the UAV's battery life is relatively inconsequential. This attribute assumes critical significance, particularly in vast spatial domains, as the allocation of the UAV's battery life to the monitoring of new informational content becomes an overriding concern. In contradistinction, our study employs real-time sensor data and progressively generates adaptable speed-based trajectories at a continuous pace over time, wherein the computational demands for online recalculations remain decoupled from the temporal prerequisites for map revisions, as they solely rely on the present image acquisition.

Stache et al. [20] proposed an IPP framework for precision agriculture, specifically targeting crop/weed segmentation, similar to our work. The distinguishing characteristic of their methodology lies in the incorporation of an accuracy model for deep learning-based architectures, enabling the quantification of the relationship between UAV altitude and semantic segmentation accuracy. They introduced a dynamic path planning approach based upon the boustrophedon method within a continuous 3D spatial domain, generating evolving trajectories at various altitudes for monitoring and close inspection tasks. Nevertheless, their approach, which involves replanning at variable altitudes, prioritizes the acquisition of higher-resolution data over minimizing flight time for comprehensive monitoring of the entire agricultural field.

One of the major factors that hinder the wider appliance of such methods is the computation needs during each replanning phase. Additionally, because several candidate paths along with their anticipated measurements should be simulated before each replanning step, their computational needs grow exponentially with respect to the field area to be covered. Previous studies have developed quite elaborate plans to mitigate this by pruning the action-space [19, 21]; however, this kind of relaxation could seriously degrade the quality of the achieved performance. Recent approaches attempt to mitigate this issue by treating the IPP as a standard Reinforcement Learning (RL) problem, learning policies that are able to compute inexpensive plans online [22]. However, the performance of this approach is highly correlated to the matching between the real world and the simulative environment with realistic data that the RL agent will be trained on. Last but not least, the majority of the available approach does not incorporate a hard constraint with respect to the available battery of the UAV, rendering their realization particularly tricky.

Aiming to overcome this "fear of missing out" important data, we propose OverFOMO, an active sensing coverage path planning approach that adopts the STC algorithm as a blueprint for the UAV path while, depending on the online received information, it adjusts its focus on specific areas. Assuming an UAV covering an agricultural field, figure 1 illustrates the proposed active sensing approach using 4 key snapshots. Snapshots 1 and 4 depict two representative examples that define the core motivation behind the proposed system, revealing that the quality of received information is inversely proportional to the UAV's speed, basically due to blurring effects. More specifically, the received image in snapshot 1 contains only a few crops and is relatively clear; therefore, the UAV can afford to speed up. Snapshot 4 presents a case where the received image is full of vegetation, but the speed of the UAV makes the segmentation process less confident. In that case (Snapshot 4), the UAV should slow down to make more accurate detections, especially in this high vegetation density subpart of the field. Snapshots 2 and 3 present the ability of the proposed active sensing scheme to handle "tricky" cases. Snapshot 2 seems to contain low vegetation coverage; however, the predicted segmentation is insecure, and therefore the UAV should speed down rapidly to verify that indeed there are no missing crops around that area. Moving to the other side of the spectrum, the received image in snapshot 3 contains much vegetation; however, the on-board segmentation process is super confident about the identification and therefore, the speed can be safely increased without sacrificing loss of information.

Although there have been proposed several alternatives to the usual practice with the back-and-forth movements [12, 13], that online calculate the next monitoring position (e.g., previously mentioned IPP methods), their time efficiency is usually significantly reduced [23, 24], leaving the back-and-forth movements as the "go-to" option for this type of missions. Within this paper, instead of proposing another approach that calculates the best next monitoring position, we strategically combine elements of the two approaches to achieve beyond state-of-the-art performance. More specifically, we keep the back-and-forth movements as the blueprint of the UAV path to also retain the performance guarantees that come with such an approach, and, at the same time, we attempt to regulate online the time spent in each sub-area based on the local information, similar to what a person would do. In a nutshell, the contributions of this work are:

- Development of a novel active sensing coverage path planning scheme that inherits the STC optimality and completeness guarantees. The computational needs for the online recalculation do not depend on the size of the operation field, making it suitable for various applications while respecting operational constraints (e.g., remaining battery, etc.).
- Development of a novel Deep-Learning-based module for adjusting the UAV speed, similar to what a human would do, taking into consideration both the quantity of the detected relevant instances (i.e. crops and weeds) and certainty (quality) about these detections. Note that the method could be extended to different operational scenarios, e.g. scan a sea area and regulate UAV speed according to marine-related classes (oil spill, algae bloom, etc.)
- An open-source, modular, simurealistic pipeline that combines the high-fidelity dynamics of AirSim [25] with real RGB images sourced from publicly available UAV datasets.
- Validation of the proposed approach, using the aforementioned simurealistic pipeline, against the widely-used STC-based coverage methods, showcasing its performance. Contrary to the prevailing state-of-the-art STC planners [8, 26], which entails a uniform scanning speed across the surveyed region, our method incorporates adaptive agent speed, resulting in reduced flight duration and enhanced image quality.

2. Problem formulation

Following the standard UAV-based monitoring Precision Agriculture (PA) process [27, 28, 29], we assume a UAV capable of acquiring images mounted with an RGB camera flying at fixed altitude. The studied problem is defined as controlling in-real-time the UAV mission parameters i.e speed, so as to acquire the best possible field representation, i.e the fidelity of field orthomosaic, within the minimum flight time.

2.1. Decision Variables

Assuming a fixed sampling rate, i.e. images per time, the IPP setup is reduced to design the series of sensing waypoints that will comprise the UAV trajectory:

$$\tau = [w_1, w_2, \ldots, w_n], \tag{1}$$

where $w_i \in \mathbb{R}^2$ denotes the image capturing position in the plane of the operational height. The time needed to complete a trajectory is denoted with $C(\tau)$ and it should be less or equal to the maximum operational flight time T_{max} of the UAV.

2.2. Field Representation Quality Assessment

After the completion of the UAV mission, all $\{I_1, I_2, \ldots, I_n\}$ images, gathered from the sampling positions as defined in (1), are going to be stitched to generate the field's orthomosaic. The quality of the extracted orthomosaic is related to the quality of the captured images $\{I_1, I_2, \ldots, I_n\}$ and proportional to the comprehensibility of the enclosed semantic content.

A common approach to measure this attribute is segmenting relative instances over the generated orthomosaic, such as crops and weeds from soil [28]. The discrimination capability of a well-trained and robust segmentation model is related to the quality of the visual input. As in the majority of semantic segmentation problems, the model efficiency gets assessed by using the Intersection over Union (IoU) [30].

2.3. Informative Path Planning Problem

Having defined the estimation approach for the field's representation quality, the general IPP problem, under the context of precision agriculture, can be translated to the following optimization problem:

$$\begin{aligned} \underset{\tau}{\text{maximize}} \quad & \frac{\mathrm{IoU}^{crop}(\tau) + \mathrm{IoU}^{weed}(\tau)}{\alpha C(\tau)} \\ \text{subject to} \quad & C(\tau) - T_{max} \leq 0 \end{aligned} \tag{2}$$

where α is used to weight $C(\tau)$ in terms of $\mathrm{IoU}^{crop}(\tau) + \mathrm{IoU}^{weed}(\tau)$, depending on the specifics of each application. For example, a usual configuration is targeting for the best possible representation (terms: $\mathrm{IoU}^{crop}(\tau) + \mathrm{IoU}^{weed}(\tau)$) within a given time budget T_{max}. For this configuration, α is chosen to be appropriate small to render the influence of the denominator technically negligible (of course, the constraint always holds).

A direct difficulty in solving (2) lies within the immense continuous domain of (1). Actually, the number of different possible combinations of (1) increases exponentially with respect to the size of the field [31], which determines the number n of image capturing positions. However, the most severe obstacle has to do with the fact that both the explicit forms of $\mathrm{IoU}^{crop}(\cdot)$ and $\mathrm{IoU}^{weed}(\cdot)$ are not available prior to the UAV mission, since ground-truth information is required. As a consequence, any approach that relies on evaluating different combinations of (1) on (2) cannot be realized within this context. On the contrary, the solution should be seeked in a method capable to assess during the ongoing mission the quality of the captured images and regulate the UAV speed accordingly, in order to extract the best possible field representation in the minimum flight time. Toward this direction, one of our main objectives is to deploy an approach that tackles (2) and its limitations, in a indirect manner.

3. Adaptive Coverage Path Planning

This section describes the details of OverFOMO, the proposed active sensing coverage path-planning algorithm, designed for previously defined optimization problem (2).

3.1. Problem Translation using Coverage Path Planning

First, let us define the Coverage Path Planning (CPP) problem [14] that is defined by the geometry of the agricultural field and the UAV characteristics. In short, Coverage Path Planning problem deals with the problem of designing a robot path that covers an area of interest in the minimum possible time. One of the most popular CPP approaches is Spanning-Tree Coverage (STC) [12] algorithm. STC first discretizes the operational area and then generates a minimum spanning-tree that will be used as a guide for the robot path. Overall, STC algorithm is a polynomial time algorithm, with respect to the field size, that guarantees complete grid coverage in the minimum possible time [12]. Hence, STC algorithm can be realized as a kernel for optimal coverage paths and generate a sequence of sensing waypoints (1), as follows:

$$x = \mathrm{STC}(\mathcal{P}, o, h, dt, s) \tag{3}$$

where $\mathcal{P}$ denotes the polygon that contains the agricultural field, o is the overlap between two images in adjacent flight path lines [32], h is the UAV flight altitude, dt denotes the time-lapse interval for the capture of each image and s is the UAV speed.

Due to the fact that we are dealing with the IPP for a specific field, $\mathcal{P}$ is considered known and constant. o and h are defined according to the specifics of each agricultural mission, e.g., plant growth rate, season, required resolution of the orthomosaic, etc. The remaining two parameters are the ones that dominate the density of the captured images (I_i from w_i position) along the STC-based path. The list of STC parameters can be further reduced, by setting the dt to its smallest feasible value for the onboard sensor that does not compromise the quality of the received images. After these realizations, STC-based trajectory for a given agricultural field and a given type of UAV can be defined as:

$$x = \mathrm{STC}(s) \tag{4}$$

Hence, utilizing (4), we now have a UAV path that completely covers the agricultural field at the minimum possible time for a given s. Inevitably, the definition of s gives rise to a trade-off. A small s value would provide premier quality on the captured images and, therefore, in our ability to distinguish accurately crops and weeds, however, it would result in covering only a small fraction of the agricultural field, due to flight time limitations. On the other hand, an increased s value could mitigate this by covering larger areas, in the expense of our discrimination accuracy. The usual practice is to apply a constant s at the beginning of the mission and perform the whole mission with such speed [8, 27, 26, 28]. However, during the operation, the UAV receives images that characterize the quantity of useful information that lies under its current path.

Within this paper, we want to exploit this online-received information and adjust the speed of the UAV during its flight, making the data acquisition process more efficient. Thus, assuming that $t_i - t_{i-1}$ denotes a fixed time-interval needed for both the information assessment and the change in the UAV speed, we want to guide the image capturing process by the following adaptive path-planning scheme:

$$x(t) = \begin{cases} \mathrm{STC}(s_1), & t_0 < t \leq t_1 \\ \mathrm{STC}(s_2), & t_1 < t \leq t_2 \\ \quad\vdots & \\ \mathrm{STC}(s_n), & t_{n-1} < t \leq C(\tau) \end{cases} \tag{5}$$

Hence, by plugging the time-varying $x(t)$ into τ, the optimization problem of (2) now is reduced to online adjust s for every time-interval of (5). In the upcoming subsections, we discuss the details of speed adjustment, i.e. calculating online $\{s_1, s_2, \ldots, s_n\}$ of (5), based on the online-received information gain at each time-interval.

3.2. Coverage Ratio & Confidence Level

Before providing the exact methodology that online adjusts the speed of the UAV, let us first define two key metrics that assess the information gain with respect to the current image frame that corresponds to the i-th time-interval in (5).

The main rationale is the fact that the information enclosed in the captured I_i image is correlated to the amount of depicted crops and weeds. Thus, a deep-learning model $\mathcal{M}$ capable of semantically segmenting images to identify three classes, namely crop, weed and background, is deployed. $\mathcal{M}$ is fed with the acquired $w \times l$ image I_i and produces a confidence score map $S_i \in \mathbb{R}^{w \times l \times 3}$ that contains the probability of each pixel belonging in each class, i.e. $S_i = \mathcal{M}(I_i)$. As a direct outcome, the prediction mask is derived using $S_i^{class} = \mathrm{argmax}\,(S_i)$, assigning a class id for every pixel. While $S_i^{prob} = \max(S_i)$ derives the overall confidence map, containing the probability of each pixel belonging to the assigned class. For improved clarity, a visual representation of the aforementioned terms is provided in figure 2.

The first metric is oriented to quantify the amount of the captured crops and weeds. To accomplish that, we utilize the *coverage ratio* (cr), inspired by [20] and defined as follows:

$$cr(S_i^{class}) = \frac{N^{crop} + N^{weed}}{N^{crop} + N^{weed} + N^{background}} \tag{6}$$

where N^{crop}, N^{weed} and $N^{background}$ denotes the number of pixels from S_i^{class} that have been classified in each class correspondingly. Note that the denominator resembles the total number of pixels in the image frame. Conceptually, $cr : \mathbb{R}^2 \rightarrow [0, 1]$ estimates the plants and weeds coverage on the target area by applying (6) rule in a pixel-wise segmented image of I_i. Low values of $cr(S_i^{class})$ imply that the vegetation enclosed in the captured image is limited and thus, the information gain of this area is low. Correspondingly, high values of cr are related to areas of lush vegetation, where the information gain is considered high.

Having this in mind, a simple formulation for the speed in (5) would be a linear mapping between cr and the speed, i.e. as the cr is increased the speed gets decreased and vice versa. However, such a formulation can have several pitfalls since the detected weed/crop instances' accuracy is not considered.

Towards this direction, the confidence of the acquired predictions is included as a second metric for assessing the information gain. More specifically, for every processed image I_i the *confidence level* (cl) is calculated as follows:

$$cl(S_i^{prob},\ S_i^{class}) = \frac{\sum_{j \in \mathcal{C}} p_j + \sum_{j \in \mathcal{W}} p_j}{N^{crop} + N^{weed}} \tag{7}$$

where $\mathcal{C}$ and $\mathcal{W}$ denote the set of pixels that have been annotated as crop and weed, respectively, and p_j denotes the corresponding confidence score for j-th pixel of S_i^{prob}.

The main idea here is that when the confidence level $cl : \mathbb{R}^2 \rightarrow [0, 1]$ of the acquired prediction is high enough, then the UAV speed can be increased to reduce the flight time since the captured image quality is adequate to make robust predictions. Respectively, a lower confidence level may imply that the quality of the processed image is low, and thus, the UAV should decrease its speed to capture a clearer view of the scene. With respect to the information gain, the confidence level can be considered as an inverse metric of the observed entropy. Higher values imply that the scene is well-known to the prediction model $\mathcal{M}$ and it can be clearly conceived; thus, the UAV can proceed faster since the acquired information is limited in this *static* environment. On the contrary, lower confidence level values imply an *unknown* environment, conceived with ambiguity; thus, speed should be decreased to increase the observation time.

3.3. Speed Adjustment

Having calculated cr (6) and cl (7) for the currently received i-th image, we can now calculate the objective speed adjustment. To perform this update we need a mapping function $G(\cdot) : \mathbb{R}^2 \rightarrow [-1, 1]$ that translates both cr and cl into speed changes with respect to the maximum allowed discrepancy q around UAV's nominal speed $\bar{s}$, i.e.

$$\begin{aligned} s_i &= \mathrm{clip}\,(u,\ \bar{s} - q,\ \bar{s} + q)\,, \\ u &= s_{i-1} + G(cr, cf)q, \text{ with } s_0 = \bar{s} \end{aligned} \tag{8}$$

where clip function constrains the updated speed between safe/acceptable bounds. Hence, to derive the needed behavior in terms of speed change, $G(\cdot)$ is defined as follows:

$$G(cr, cf) = \omega_1(cl)g_1(cr) + \omega_2(cl)g_2(cl) \tag{9}$$

where $g_1(\cdot)$ and $g_2(\cdot)$ denote the translation functions from cr and cl, respectively, to a relative speed change. Additionally, for each term a regulation function is defined, namely $\omega_1(\cdot)$ and $\omega_2(\cdot)$, to prioritize one term over the other. $g_1(\cdot)$ and $g_2(\cdot)$ have chosen to be linear piecewise functions, while $\omega_1(\cdot)$ and $\omega_2(\cdot)$ are of type of parabola with respect to their parameters. For ease of understanding, figure 3 graphically illustrates the form of these functions. Additional information regarding the calibration of $g(\cdot)$ and $w(\cdot)$ functions is provided in Appendix B.

Note that the speed adjustment s_i in position w_i is with respect to the previous speed "state" s_{i-1}, instead of the nominal speed

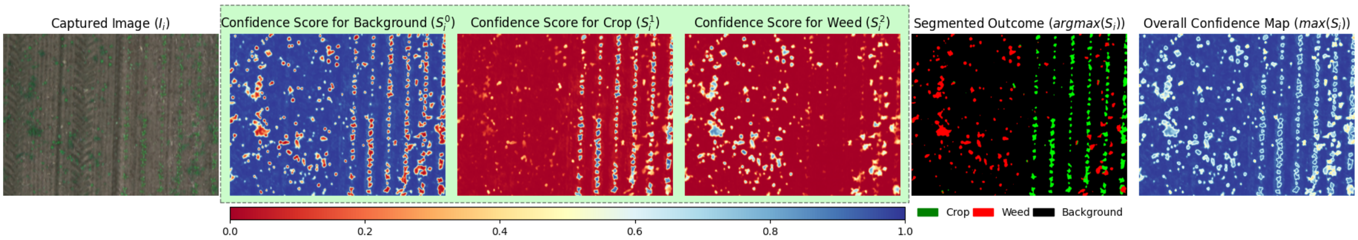

Figure 2: Demonstration example of the semantic segmentation model output which is employed to calculate cr and cl metrics. Captured image I_i is fed to the model and produces the 3-channel array S_i (illustrated per channel and highlighted with pale-green color), where each channel contains the probability of the image pixels belonging to the corresponding class. $S_i^{class} = \text{argmax}\,(S_i)$ leads to the segmented outcome, where each pixel is assigned to one of the 3 classes, enabling the estimation of cr metric. The overall confidence map, providing the probability of each pixel belonging to the assigned class, is acquired via $S_i^{prob} = \max(S_i)$ and enables the calculation of cl metric.

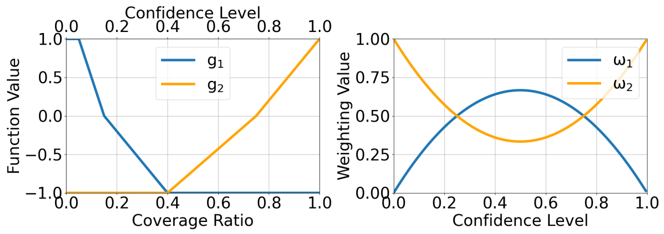

Figure 3: Graphical illustration of the employed functions in (9). Translation functions (left) $g_1(\cdot)$ and $g_2(\cdot)$ aim to map the calculated cr and cl, respectively, to a relative speed change. Weighting functions (right) $\omega_1(\cdot)$ and $\omega_2(\cdot)$ aim to regularize the contribution of $g_1(\cdot)$ and $g_2(\cdot)$ to the final decision.

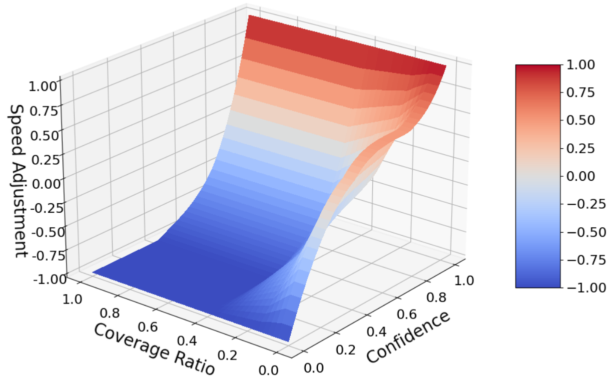

Figure 4: 3D graph of the designed $G(\cdot)$ function to adapt UAV speed according to the information gain.

$\bar{s}$. The specific choice enables more smooth transitions of the vehicle speed, while the adapting process can be considered to some extent stateful. Furthermore, contrary to the established approaches, the presented method does not adjust the overlap among consecutively captured images to a fixed value [29]. To this end, tuning parameter q is enabled to regulate the range of the speed adjustments and, thus, maintain the image overlap within acceptable (application-wise) thresholds [33]. Towards this direction, the proposed method aims to control the quality of the captured image data by adjusting the vehicle speed (with respect to the semantic content of the scene) and, thus, regulating the image distortion due to motion blurring. Both q and $\bar{s}$ are user-defined parameters that can express both the user requirement and the UAV hardware characteristics.

The main rationale of weighting functions $\omega_1(\cdot)$ and $\omega_2(\cdot)$ in (9) is to adjust the contribution of each term based on the confidence of the prediction. For instance, assuming that cl value is 0, then the estimated value of cr and, by extension the value of $g_1(cr)$ is irrelevant since it is based on inaccurate predictions. Similarly, in the case of $cl = 0.5$ the prediction can be considered to some extent as ambiguous and the formulation favors coverage ratio measurements in order to regulate speed[1]. Aiming to provide further insights regarding the system's behavior under different scenarios, in figure 4 is demonstrated a 3D representation of $G(\cdot)$ function for its whole domain.

[1]Please note that, although the information gain can be described quite effectively by these functions, their forms can be further fine-tuned to achieve better, problem-oriented performance.

3.4. Proposed Method as a Whole

Having analyzed the key points in the previous sections, the proposed adaptive path planning can be summarized as "estimate the information gain captured in image I_i and adapt the vehicle speed according to it". Since there is no ground truth, we employ the two metrics, coverage ratio (cr) and confidence level (cl), in order to tackle its absence and concurrently quantify the information enclosed in each image. Coverage ratio estimates the amount of crops and weeds in the scene and aims to answer the question "how much significant is this area?". Confidence level aims to quantify the validity of the model estimation and responds to the question "how much accurate though is the estimation regarding the significance of this specific area?".

At each step, the proposed method answers these two questions and regulates the UAV speed accordingly through function $G(\cdot)$. In figure 5 we present a comprehensive set of operational scenarios, providing insights regarding the expected behavior of an adaptive system that self-regulates its speed, which was our main motivation, along with the key-values of the Over-FOMO that lead to the corresponding adjustment. The first two rows of the figure refer to cases where the model can provide a concrete estimation regarding the amount of existing crops or weeds and the UAV speed is regulated according to the quantity

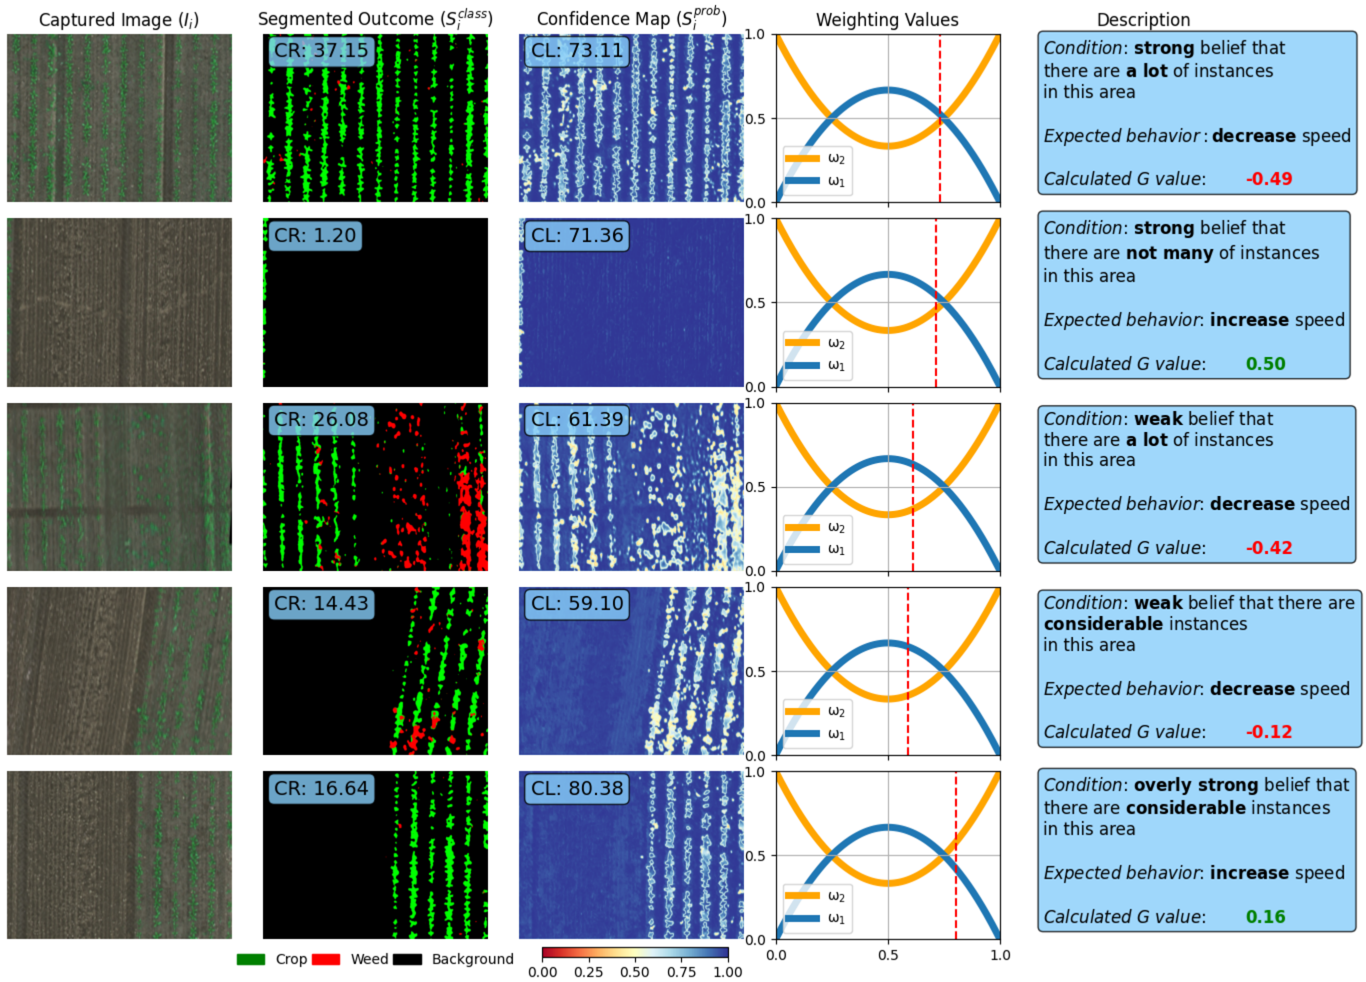

Figure 5: Illustration of different operational cases of the adaptive coverage path planning. Each row presents the analysis conducted for the corresponding captured image. For each case, the corresponding *cr* and *cl* metrics are mentioned (%). Vertical red line in the figures of the fourth column corresponds to the *cl* metric, based on which the weighting values are calculated and employed in (9) to calculate the corresponding G value. Last column provides a short description of each case among with the expected behavior of an adaptive system and the corresponding G value of our method, which is employed in (8) to update the UAV speed accordingly.

of the detected instances. Rows 3 and 4 refer to cases where the quality of captured data deteriorated due to motion blurring. One can notice the impact of this effect on the calculated cl metric. Despite the amount of estimated crops and weeds, the vehicle speed is decreased since the quality of captured data implies ambiguous estimations. The last row resembles the case where the estimator is overly confident implying that the data quality is adequate and therefore a partial deterioration, by increasing the UAV speed, can be tolerated to save flight time. As demonstrated, the proposed adaptive scheme can confront variable cases. In this direction, the proposed adaptive scheme considers the quantity (cr) and the quality (cl) of the information gained per image, aiming to operate in a sweet spot where the quality of captured data is maximized while the flight time is minimized.

In a nutshell, Algorithm 1 outlines the proposed adaptive coverage path planning as a whole. Putting everything together, the proposed approach alleviates both the combinatory nature and the unknown factor by a careful combination of two ingredients: i) the STC algorithm that is capable of computing offline optimal coverage paths with $O(n)$ complexity, and ii) an online speed adjustment scheme that takes into consideration the current information gain.

Algorithm 1 Adaptive Coverage Path Planning

Require: $\mathcal{P}, o, h, dt, \bar{s}, q, \mathcal{M}$
Ensure: τ

Offline phase:
Define a STC-based trajectory parametric over s (5)
Online phase:
for each viewpoint w_i at t_i **do**
  Acquire frame I_i
  $S_i \leftarrow \mathcal{M}(I_i)$ and $S_i^{class} \leftarrow \text{argmax}\,(S_i)$
  Calculate cr and cl according to (6) and (7)
  $G(cr, cf) = \omega_1(cr)g_1(cr) + \omega_2(cl)g_2(cl)$
  $s_i \leftarrow$ apply (8)
end for

4. Experimental Evaluation

In this section, our active sensing planning approach is evaluated via a simu-realistic pipeline by incorporating a high-fidelity simulator and a large-scale dataset for precision agriculture applications.

4.1. Dataset

The exploited dataset was WeedMap [27], which contains multi-spectral images from sugar beet crops and weeds interfering in the crop lines. Data were collected during two campaigns, the first led to 3 orthomosaic maps while the second to 5. For every map the depicted plants were pixel-wise annotated, leading to 3 different classes, namely crop, weed and background. Every orthomosaic is provided also in a tiled version, where the original image is divided into patches of 480×360 pixels. In our case, only RGB data from the second campaign were utilized.

4.2. Detection Model

Regarding the detection model $\mathcal{M}$ that semantically segments crop and weed instances, a deep-learning method was utilized. In specific, the well-known UNet [34] architecture enhanced with EfficienNetB1 [35] network as backbone was employed. This design was selected based on the balanced trade-off amongst inference time and model accuracy, taking into consideration that our aim was to deploy a real-time operating system. The deployed model was trained on WeedMap for 500 epochs. A set of image processing techniques was utilized for data augmentation, in specific, image rotation, resize, vertical/horizontal flip and brightness change. At last, 255×255 patches were randomly cropped from the tiled input images. Training was conducted with Adam optimizer with learning rate and batch size equal to 10^{-3} and 16, respectively.

4.3. Setup

To evaluate the proposed method in the context of the aforementioned dataset and assess its performance as a real-time interaction system, a hybrid simu-realistic framework was designed. The main goal here is to simulate real-world missions with real-time interactions, in order to generate the required set of viewpoints w_i, collect the corresponding images I_i and produce the most optimal field representation i.e a 2D orthomosaic within the minimum operational time.

To accurately simulate the UAV's physics and dynamics and emulate its motion control, AirSim [25], an open-source high-fidelity simulator for autonomous vehicles, was utilized. AirSim is capable of forwarding the world dynamics, including a wide range of weather dynamics, at a high frequency allowing for real-time, hardware-in-the-loop ready, realistic simulations. All the experiments were carried out with a single drone within AirSim platform. The geo-referenced orthomosaic images of WeedMap fields, allow the direct mapping of the simulated UAV location in world coordinates to the pixel-level coordinates of the corresponding field's orthophoto. Thus, the exploited testing fields can be considered as natural parts of the environment and the UAV's camera input can be simulated by cropping image patches from the related orthophoto, with respect to the vehicle position. In Table 1 are provided further details regarding the parameters related to the UAV flight and the simulated camera sensor. The selection was based on the corresponding information provided in WeedMap dataset.

Table 1: Path planning and sensor specifications.

Type	Description	Specification	Unit
UAV System	Flight altitude	10	meters
	f_{best}	1	frame/sec
Visual Sensor	Overlap	70	%
	Gimbal pitch	-90	degrees
	Image size (width × height)	640×480	pixels

In this light, for each agricultural field in the deployed dataset, QGIS platform[2] was used to specify the filled-in polygon $\mathcal{P}$ of (3) and an STC-based coverage path was designed and integrated into AirSim. During the simulated flight with initial speed $\bar{s}$, a set of processing operations are applied in a recursive manner in order to adapt the UAV speed in real-time. The core loop of this process is illustrated in figure 6. In specific, with time interval f_{best}, the coordinates of drone viewpoint w_i are extracted from AirSim environment. The acquired point is mapped to the corresponding geo-referenced orthomosaic image of the examined field and a 640×480 image is cropped according to the AirSim-emulated UAV trajectory. Furthermore, motion blur is applied to the cropped image according to the current UAV speed, aiming to create realistic captured data. In specific, we followed the formulation presented in [36]. Assuming there is no additive noise, the blurred image B_i is simply acquired by the convolution of a blur kernel K with the captured image I_i, i.e $B_i = K * I_i$. We know that the UAV is moving in the same direction as the vertical axis of the captured images. Thus, the blur kernel K can be easily emulated with a vertical kernel (ones in the middle column and zeros everywhere else). In order to simulate the blurring effect impact according to vehicle speed, we increased the kernel size, e.g. $3 \times 3, 5 \times 5$, etc, respectively. Next, the acquired image I_i is forwarded to the proposed adaptive scheme, that assess the information gain enclosed in it and adapts the vehicle speed based on the proposed translation function $G(\cdot)$. The update information is fed back to the simu-realistic environment, regulating the UAV speed on-the-fly. The aforementioned process is repeated at the next time interval, for viewpoint w_{i+1}.

4.4. Baseline

To evaluate the efficiency of the proposed speed adjustment methodology, we chose to compare it against the "go-to" STC-based coverage path-planning approach for precision agriculture applications [14, 8, 26], where the UAV is moving with constant speed. Note that the flight path in both scenarios is identical, while the sampling interval remains the same in all cases. However, variations in speed lead to collecting data from different viewpoints w_i. Moreover, according to the vehicle speed during

[2] https://qgis.org/en/site/

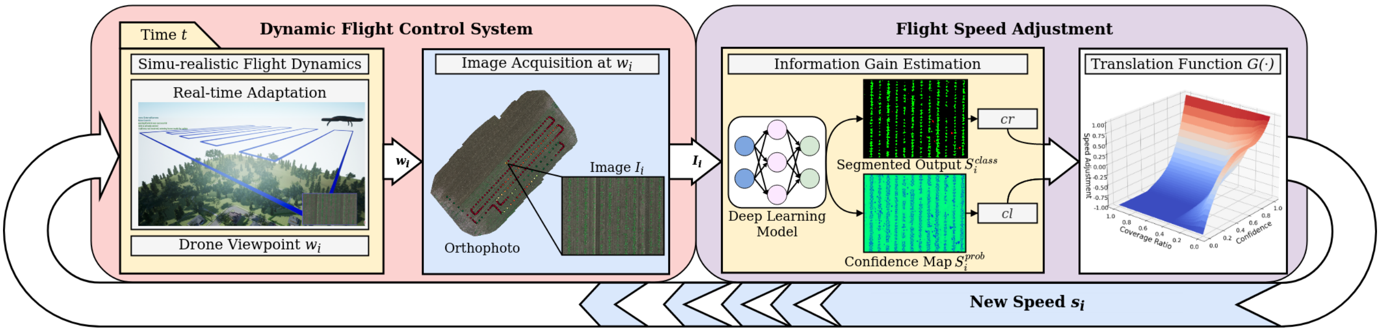

Figure 6: Graphical illustration demonstrating the core loop of the designed experimental setup. A simu-realistic flight environment, based on AirSim, is deployed to produce, in real time, UAV viewpoint w_i. Image I_i is cropped at w_i position from the field orthophoto and blurred according to current UAV speed. I_i is processed by the proposed adaptive scheme to estimate the information gain of the scene and adjust UAV speed to s_i, based on the designed $G(\cdot)$ translation function.

the capturing time, acquired images differ in terms of image quality due to motion blurring. We refer to the deployed non-adaptive method as *STC-PA*, while the proposed method is mentioned as *OverFOMO*. Through this comparison, we aim to answer the following question: instead of covering the field with constant speed $\bar{s}$, can the speed adjustments of the proposed method lead to more meaningful data in less or comparative time?

4.5. Performance Analysis

The proposed method was extensively evaluated under different flight scenarios and agricultural environments. More specifically, for each one of the 5 crop areas, we deployed the adaptive planning process through the aforementioned simurealistic pipeline, for different selections of nominal speed, in m/s, namely $\bar{s} \in \{3, 4, 5, 6\}$. q parameter of (8) was set to 1 implying that UAV can increase or decrease its nominal speed by 1 m/s at maximum. The evaluation process is based on recreating the orthomosaic map from the set of images I collected during the *OverFOMO* mission. Next, the stitched outcome is semantically segmented, utilizing the aforementioned trained model, and IoU is calculated for crop and weed class. Our aim is to quantify the quality of the reconstructed map in terms of the enclosed semantic content and thus, provide a metric of the scanning efficiency of the planned mission.

The aforementioned evaluation process is applied for each one of the examined fields, computing the execution time and the IoU for crop and weed class. In order to conduct credible validations, in each case the testing field is excluded from the training process of the detection model. The same approach is followed for both the *STC-PA* and the *OverFOMO* approach. The two methods are compared in figure 7, where is presented the average IoU over the 5 examined fields and its variance for crop and weed class correspondingly, for different nominal speeds. In total, 20 flight scenarios (4 nominal speeds × 5 fields) were executed for each of the two evaluated approaches. Furthermore, for the *STC-PA* method we examine the case of $\bar{s} = 2\ m/s$ which is considered as the ideal scenario, where the UAV is moving with the minimum speed and thus, data are collected totally undistorted (no motion blur is applied).

For both classes, the proposed method outperforms *STC-PA*. The efficiency of adaptive planning, in terms of IoU, is clear in case of crop detection, while in case of weed the maximum IoU of the proposed method is constantly higher than the comparative for the whole set of examined nominal speeds. In terms of execution time, for lower values of nominal speed, the adaptive method is to some extent slower yet, in favor of higher accuracy. As the nominal speed increases the execution time gap between the two methods is decreased, while for $\bar{s} = 6\ m/s$, the proposed method outperforms the *STC-PA* in terms of flight time also. All in all, results imply that *OverFOMO* scans efficiently an examined area, collecting high quality data from the areas containing rich semantic content while passing by areas of lower interest to reduce the flight time.

4.6. Qualitative Analysis

In order to validate further the efficiency of the developed method we evaluate the generated orthomosaic maps in terms of image quality. Towards this direction, simulated missions deployed with the *STC-PA* and the *OverFOMO* method are conducted for the field “002” of the Weedmap dataset, with nominal speed $\bar{s} = 3\ m/s$. Next, we estimate the image similarity, in terms of Structural Similarity Index (SSIM) [37], among the original orthomosaic (provided in the dataset) and the one built via data collected from the *OverFOMO* mission. The same process is followed for the *STC-PA* method. In figure 8 qualitative results for the two comparative approaches are presented. More specifically, in figure 8(a) the original orthomasaic image is presented, while in figure 8(b) is illustrated the annotated ground truth, aiming to provided further insights regarding the semantic content of the examined scene. In figure 8(c) & (d) the estimated SSIM index is demonstrated for the cases of *STC-PA* and *OverFOMO*, respectively. For visualization purposes, the similarity of the generated orthomosaics to the original one is illustrated in red-blue colorscale. Blue areas indicate higher similarity, while yellow and red regions indicate deviations between the generated and the original image. For a more comprehensive comparison, the histogram of the calculated SSIM values is provided for each case in figure 8(e).

Results imply that the proposed method leads to a more accurate orthomosaic map compared to the current “go-to” approach, especially for areas of high information gain, where the measured similarity is higher (note dark blue regions in figure 8(d)). The fidelity of the generated orthomosaic indicates that the collected data of the adaptive mission can enclose more precisely

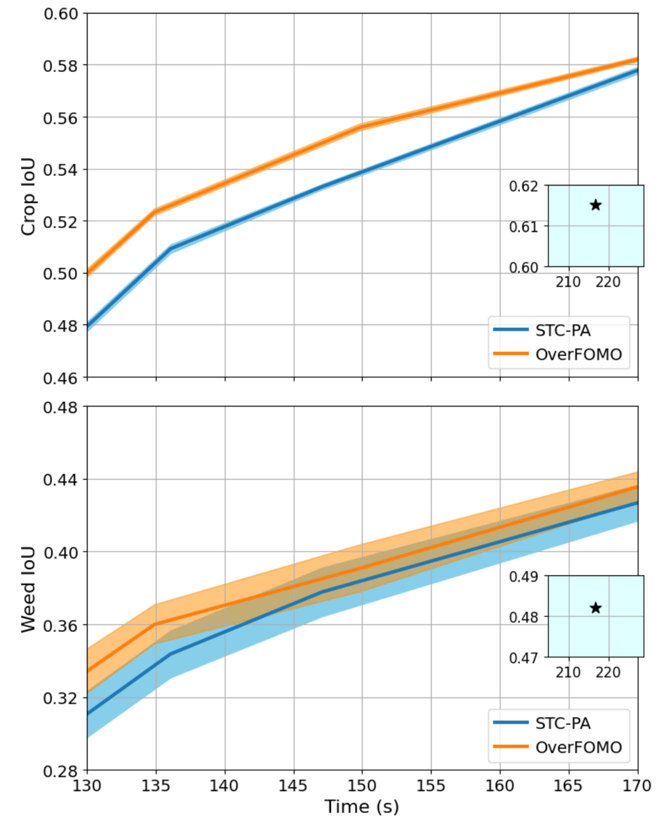

Figure 7: Averaged IoU for the testing fields of Weedmap dataset. Solid line refers to mean value and shaded region to variance. Black star refers to the ideal scenario where *STC-PA* method is applied with nominal speed $\bar{s} = 2\ m/s$ and can be considered as the convergence point of the two methods. For both crop (top) and weed (bottom) classes, the proposed adaptive method leads to higher performance, implying more accurate scanning of the examined area.

the semantic content of the scene. This is also supported by the provided histograms in figure 8(e), where the proposed method reports higher similarity values for the majority of cases. Taking into consideration the overall patterns of SSIM index values, with respect to the information of figure 8(b), one can derive that the proposed method regulates the vehicle speed according to the semantic content of the scanned field. In areas where the information gain is high, i.e. lush vegetation, UAV decelerates to acquire high quality - less blurry - data, while in areas of lower interest it accelerates since the information gain is considered minimum. On the contrary, the *STC-PA* method of constant speed scanning presents, to some extent, constant image quality levels, distributed across the whole field, without taking into consideration the semantic content of the scene.

5. Conclusions

In this work an UAV active sensing coverage path planning scheme for precision agriculture tasks has been presented. Our method is capable of adjusting the UAV speed based on the perceived visual information (i.e. observed crops and weeds), while the computational needs for the online processing are uncoupled to the operation field's size. A core-element of the proposed approach is a robust deep learning-based module, allowing to regulate the vehicle speed according to the quantity of the detected instances and the quality (confidence) of such detections. The proposed method has been extensively validated through a designed simu-realistic environment, conducting several missions with different nominal speed for 5 different agricultural fields of WeedMap dataset. Compared to the well-known lawnmover coverage path planning, our method manages to capture higher quality data in comparable execution times. In the future we aim to deploy our method in real-world scenarios by employing UAVs with on-board capabilities.

6. Acknowledgment

This research has been financed by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH - CREATE - INNOVATE (T1EDK-00636). We gratefully acknowledge the support of NVIDIA Corporation with the donation of GPUs used for this research.

Appendix A. Image Quality vs Method Performance

In this appendix are presented further details regarding how the proposed method's performance is affected from the efficiency of the employed $\mathcal{M}$ model, the deduction of image quality due to speed increment and the possible misclassifications.

More specifically, the clarity of the on-the-fly captured images affects the segmentation confidence during the online phase of the adaptive system. Through the extensive evaluation of the deployed deep-learning model, we noticed that motion blur mostly affects the clearness of the depicted weeds and crops, increasing the ambiguity of their exact shape and size, and under the perspective of Bayesian modeling [38], increasing the *aleatoric* uncertainty. *Epistemic* uncertainty is also inherent in the prediction system, and it is reflected in the deviation of the captured image from the distribution of the training data. The offline validation of the employed semantic segmentation model, implied that it can generalize well in previously unseen data and thus, the effect of the epistemic uncertainty is not crucial. However, training data refer to an ideal scenario where the utilized images contain no distortion. Thus, during the online phase, aleatoric uncertainty expressed through the blurring effect significantly affects the efficiency of the on-the-fly prediction.

The above analysis comprises the challenging nature of the problem that we aim to tackle. Towards this direction, we use this uncertainty to our advantage in order to regulate the UAV speed according to it. By considering the confidence score of the segmented outcome, through the *cl* metric, we aim to estimate the information gain at each sensing waypoint in respect to the confidence of this estimation.

In figure A.9 we present the proposed method performance for an input image which is gradually deteriorated via motion blurring. One can notice that although the *cr* metric is slightly decreased, the *cl* value is significantly dropped, implying uncertainty in the acquired estimations and deterioration of the image

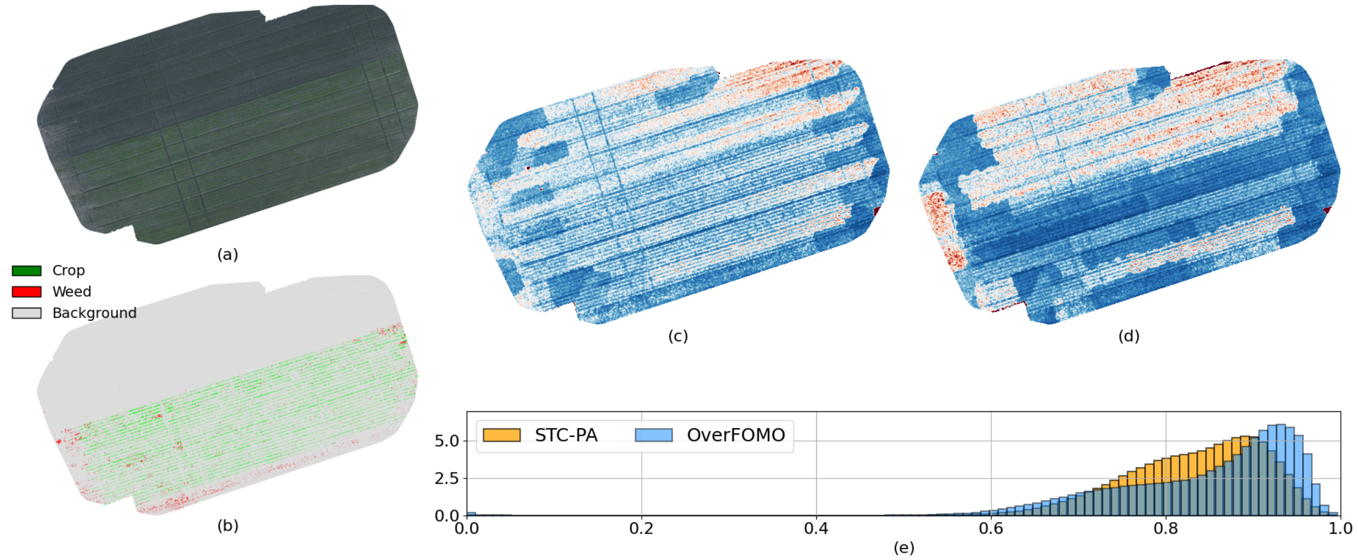

Figure 8: Qualitative results for the *STC-PA* and the proposed *OverFOMO* approach. In (a) is presented the original orthomosaic image, while in (b) the semantic content of the examined scene. In (c) and (d) is illustrated the image similarity, in terms of SSIM, among the original (b) and the generated orthomosaic via coverage missions planned following the *STC-PA* and the *OverFOMO* method, respectively. Both missions are conducted with same nominal speed. Red colors resemble lower values of SSIM, while higher values of SSIM are mapped with blue colors.In (e) is presented the corresponding histogram of the calculated SSIM index for both cases.

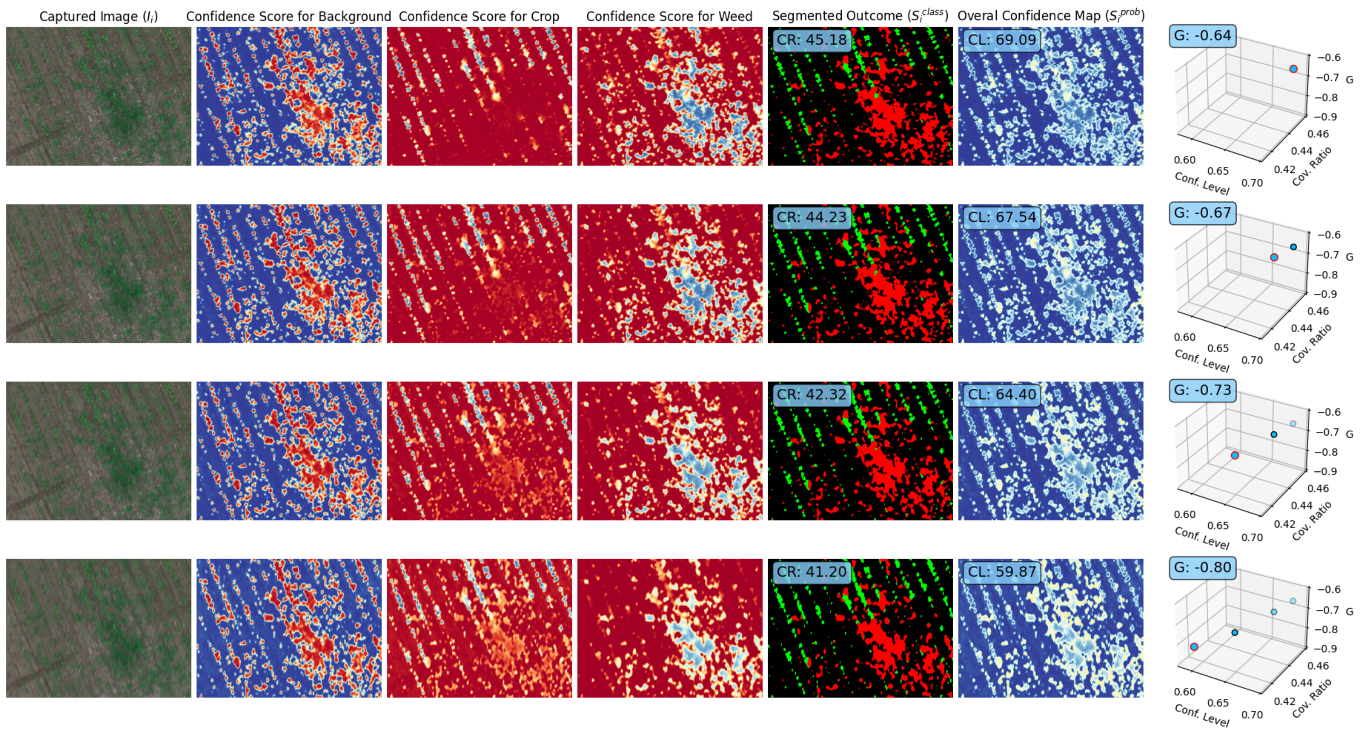

Figure A.9: Illustration of the impact of motion blur on the proposed method performance. In each row, the motion blurring applied to the input image is increased, leading to lower values of cl (%) metric, although cr (%) remains at similar levels. Last column presents in the $3D$ space the calculated G value (red outline) among with the corresponding values of the previous blurring cases. One can note the gradual decrease of G value, implying the reduction of vehicle speed. All in all, the proposed OverFOMO approach takes into consideration the confidence of the segmentation model and adjusts the UAV speed accordingly to acquire more accurate estimations that meet the application-oriented requirements.

quality due to the enhancement of the blurring effect. Please note the calculated G values in all cases, which are gradually decreased, implying the adjustment of speed to lower values. The presented illustration demonstrates the ability of the proposed method to adapt the vehicle speed in order to avoid missing vital information and cope with possible misclassifications due to low image quality.

Appendix B. Calibration of Translation and Weighting Functions

In order to obtain the $g(\cdot)$ and $w(\cdot)$ functions of figure 3 we followed a reverse engineering approach to make the adaptive system meet the expected behavior of the characteristic cases presented in figure 5. According to the presented formulation for $G(\cdot)$, we want the translation functions $g_1(\cdot), g_2(\cdot)$ to map the input to values from -1 to 1. Similarly, the weight functions $w_1(\cdot), w_2(\cdot)$ should range from 0 to 1 and sum to 1. Based on that, we focused on a family of linear-wise and parabola functions for $g(\cdot)$ and $w(\cdot)$, respectively. Moreover, during the training and evaluation of the $\mathcal{M}$ model that segments the images, we acquired valuable insights. First, we know that since it is a 3 class problem, the probability p_j cannot be lower than 0.33. Thus, we do not expect values lower than that for cl metric. Moreover, by examining sample images of the evaluation set we concluded that adequately accurate detections are acquired when cl is around 0.75, thus we considered this as a break-point. Similarly, we noticed that at the current altitude the peak coverage ratio is around 0.4 while cr values below 0.15 refer to areas of low vegetation. Regarding the $w(\cdot)$ parabola functions, they were designed to control the contribution of each metric and express the system's expected behavior. We want to ignore the estimated coverage ratio in case that this estimation is ambiguous or overly strong. In case of moderate belief, we want the system to be guided accordingly, taking also into consideration the cr value. Please note that the presented functions are not the unique solution, even for the specific IPP problem. One can select different functions, or tune their key-points according to the use-case and in respect to how much tolerance can be enclosed to the information quality - speed trade off.

References

[1] I. Martinez-Alpiste, G. Golcarenarenji, Q. Wang, J. M. Alcaraz-Calero, Search and rescue operation using uavs: A case study, Expert Systems with Applications 178 (2021) 114937.

[2] D. Kim, M. Liu, S. Lee, V. R. Kamat, Remote proximity monitoring between mobile construction resources using camera-mounted uavs, Automation in Construction 99 (2019) 168–182.

[3] H. X. Pham, H. M. La, D. Feil-Seifer, M. C. Deans, A distributed control framework of multiple unmanned aerial vehicles for dynamic wildfire tracking, IEEE Transactions on Systems, Man, and Cybernetics: Systems 50 (2018) 1537–1548.

[4] J. Rodríguez, I. Lizarazo, F. Prieto, V. Angulo-Morales, Assessment of potato late blight from uav-based multispectral imagery, Computers and Electronics in Agriculture 184 (2021) 106061.

[5] N. O'Mahony, S. Campbell, A. Carvalho, S. Harapanahalli, G. V. Hernandez, L. Krpalkova, D. Riordan, J. Walsh, Deep learning vs. traditional computer vision, in: Science and information conference, Springer, 2019, pp. 128–144.

[6] J. González-García, R. L. Swenson, A. Gómez-Espinosa, Real-time kinematics applied at unmanned aerial vehicles positioning for orthophotography in precision agriculture, Computers and Electronics in Agriculture 177 (2020) 105695.

[7] Y. Ampatzidis, V. Partel, L. Costa, Agroview: Cloud-based application to process, analyze and visualize uav-collected data for precision agriculture applications utilizing artificial intelligence, Computers and Electronics in Agriculture 174 (2020) 105457.

[8] G. D. Karatzinis, S. D. Apostolidis, A. C. Kapoutsis, L. Panagiotopoulou, Y. S. Boutalis, E. B. Kosmatopoulos, Towards an integrated low-cost agricultural monitoring system with unmanned aircraft system, in: 2020 International conference on unmanned aircraft systems (ICUAS), IEEE, 2020, pp. 1131–1138.

[9] M. D. Bah, E. Dericquebourg, A. Hafiane, R. Canals, Deep learning based classification system for identifying weeds using high-resolution uav imagery, in: Science and Information Conference, Springer, 2018, pp. 176–187.

[10] J. Rückin, L. Jin, F. Magistri, C. Stachniss, M. Popović, Informative path planning for active learning in aerial semantic mapping, arXiv preprint arXiv:2203.01652 (2022).

[11] A. Meliou, A. Krause, C. Guestrin, J. M. Hellerstein, Nonmyopic informative path planning in spatio-temporal models, in: AAAI, volume 10, 2007, pp. 16–7.

[12] Y. Gabriely, E. Rimon, Spanning-tree based coverage of continuous areas by a mobile robot, Annals of mathematics and artificial intelligence 31 (2001) 77–98.

[13] H. Choset, P. Pignon, Coverage path planning: The boustrophedon cellular decomposition, in: Field and service robotics, Springer, 1998, pp. 203–209.

[14] T. M. Cabreira, L. B. Brisolara, F. J. Paulo R, Survey on coverage path planning with unmanned aerial vehicles, Drones 3 (2019) 4.

[15] C. K. Williams, C. E. Rasmussen, Gaussian processes for machine learning, volume 2, MIT press Cambridge, MA, 2006.

[16] G. Hitz, E. Galceran, M.-È. Garneau, F. Pomerleau, R. Siegwart, Adaptive continuous-space informative path planning for online environmental monitoring, Journal of Field Robotics 34 (2017) 1427–1449.

[17] K. C. Vivaldini, T. H. Martinelli, V. C. Guizilini, J. R. Souza, M. D. Oliveira, F. T. Ramos, D. F. Wolf, Uav route planning for active disease classification, Autonomous robots 43 (2019) 1137–1153.

[18] M. Popović, T. Vidal-Calleja, G. Hitz, I. Sa, R. Siegwart, J. Nieto, Multiresolution mapping and informative path planning for uav-based terrain monitoring, in: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2017, pp. 1382–1388.

[19] M. Popović, T. Vidal-Calleja, J. J. Chung, J. Nieto, R. Siegwart, Informative path planning for active field mapping under localization uncertainty, in: 2020 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2020, pp. 10751–10757.

[20] F. Stache, J. Westheider, F. Magistri, C. Stachniss, M. Popović, Adaptive path planning for uavs for multi-resolution semantic segmentation, Robotics and Autonomous Systems 159 (2023) 104288. URL: https://www.sciencedirect.com/science/article/pii/S0921889022001774. doi:https://doi.org/10.1016/j.robot.2022.104288.

[21] D. I. Koutras, A. C. Kapoutsis, E. B. Kosmatopoulos, Autonomous and cooperative design of the monitor positions for a team of uavs to maximize the quantity and quality of detected objects, IEEE Robotics and Automation Letters 5 (2020) 4986–4993.

[22] J. Rückin, L. Jin, M. Popović, Adaptive informative path planning using deep reinforcement learning for uav-based active sensing, in: 2022 International Conference on Robotics and Automation (ICRA), IEEE, 2022, pp. 4473–4479.

[23] A. C. Kapoutsis, S. A. Chatzichristofis, L. Doitsidis, J. B. de Sousa, J. Pinto, J. Braga, E. B. Kosmatopoulos, Real-time adaptive multi-robot exploration with application to underwater map construction, Autonomous robots 40 (2016) 987–1015.

[24] A. Renzaglia, L. Doitsidis, A. Martinelli, E. B. Kosmatopoulos, Multi-robot three-dimensional coverage of unknown areas, The International Journal of Robotics Research 31 (2012) 738–752.

[25] S. Shah, D. Dey, C. Lovett, A. Kapoor, Airsim: High-fidelity visual and physical simulation for autonomous vehicles, in: Field and service robotics, Springer, 2018, pp. 621–635.

[26] T. H. Pham, Y. Bestaoui, S. Mammar, Aerial robot coverage path planning approach with concave obstacles in precision agriculture, in: 2017 Workshop on Research, Education and Development of Unmanned Aerial Systems (RED-UAS), IEEE, 2017, pp. 43–48.

[27] I. Sa, M. Popović, R. Khanna, Z. Chen, P. Lottes, F. Liebisch, J. Nieto, C. Stachniss, A. Walter, R. Siegwart, Weedmap: A large-scale semantic weed mapping framework using aerial multispectral imaging and deep neural network for precision farming, Remote Sensing 10 (2018) 1423.

[28] M. Krestenitis, E. K. Raptis, A. C. Kapoutsis, K. Ioannidis, E. B. Kosmatopoulos, S. Vrochidis, I. Kompatsiaris, Cofly-weeddb: A uav image dataset for weed detection and species identification, Data in Brief 45

(2022) 108575.

[29] P. Radoglou-Grammatikis, P. Sarigiannidis, T. Lagkas, I. Moscholios, A compilation of uav applications for precision agriculture, Computer Networks 172 (2020) 107148.

[30] H. Rezatofighi, N. Tsoi, J. Gwak, A. Sadeghian, I. Reid, S. Savarese, Generalized intersection over union: A metric and a loss for bounding box regression, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 658–666.

[31] E. Galceran, M. Carreras, A survey on coverage path planning for robotics, Robotics and Autonomous systems 61 (2013) 1258–1276.

[32] S. D. Apostolidis, P. C. Kapoutsis, A. C. Kapoutsis, E. B. Kosmatopoulos, Cooperative multi-uav coverage mission planning platform for remote sensing applications, Autonomous Robots (2022) 1–28.

[33] D. C. Tsouros, S. Bibi, P. G. Sarigiannidis, A review on uav-based applications for precision agriculture, Information 10 (2019) 349.

[34] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer, 2015, pp. 234–241.

[35] M. Tan, Q. Le, Efficientnet: Rethinking model scaling for convolutional neural networks, in: International conference on machine learning, PMLR, 2019, pp. 6105–6114.

[36] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, W. T. Freeman, Removing camera shake from a single photograph, in: Acm Siggraph 2006 Papers, 2006, pp. 787–794.

[37] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE transactions on image processing 13 (2004) 600–612.

[38] A. Kendall, Y. Gal, What uncertainties do we need in bayesian deep learning for computer vision?, Advances in neural information processing systems 30 (2017).