2022 International Conference on Unmanned Aircraft Systems (ICUAS)
Dubrovnik, Croatia. June 21-24, 2022

2022 International Conference on Unmanned Aircraft Systems (ICUAS) | 978-1-6654-0593-5/22/$31.00 ©2022 IEEE | DOI: 10.1109/ICUAS54217.2022.9836120

Deep-Learned Autonomous Landing Site Discovery for a Tiltrotor Micro Aerial Vehicle

Prateek Arora*, Stephen J. Carlson, Tolga Karakurt, Christos Papachristos

***Abstract*— This paper addresses the need for autonomous Micro Aerial Vehicle (MAV) landing in unstructured environments. For aerial robotics to become truly ubiquitous, the requirement for operation over generic environments where a priori knowledge of landing site locations is unavailable has to be met. We propose the utilization of a Deep Learning-based classification framework deployed on aerial view image captured by the aerial robot, to perform the functionalities of object detection and classification below the aircraft, in order to autonomously characterize potential sites that are safe for landing. The proposed solution employs a small-sized companion computer equipped with a Neural Processing Unit, capable of handling both the high-level autonomy tasks, as well as executing the Neural Network pipeline onboard a MAV–class aircraft. We demonstrate the system's effectiveness by deploying it onboard the MiniHawk-VTOL, a custom-developed hybrid flight envelope aerial robot with solar energy harvesting capabilities, designed to accommodate the aforementioned needs with the ultimate goal of enabling autonomous migratory missions (via anytime/anyplace land-to-recharge over unknown, unstructured environments).**

I. INTRODUCTION

Autonomy in aerial robotics remains a rising trend and an increasingly pursued area of research over the past decade. With vast real world utility, unmanned aerial systems are being deployed in retail & e-commerce, search and rescue [1–4], industrial inspection [5–10], and exploration of both terrestrial challenging environments [11–17] and of other planets as well [18, 19]. In becoming increasingly ubiquitous, a core requirement is to perform autonomous missions beyond visual line of sight and unattended.

"Migratory" Unmanned Aerial Vehicles (UAVs) [20] can be defined as a class of such systems capable of performing long–distance missions without any human assistance, via landing-to-recharge their battery leveraging in situ energy harvesting capabilities. This can be facilitated by GPS-based solutions for self-localization and landing at predetermined georeferenced locations, which allow for safe landing and recharging. However, for a truly autonomous UAV characterizing and georeferencing every possible landmark along any arbitrary path would be impossible. Additionally, any possible GPS malfunction or degradation, or even adversarial attacks, should not compromise its capacity to at least avoid no-fly zones above critical structures. To this purpose, an intelligent solution for identifying significant terrain landmarks (for landing or avoidance) relying exclusively on onboard real-time aerial imagery is considered within this work.

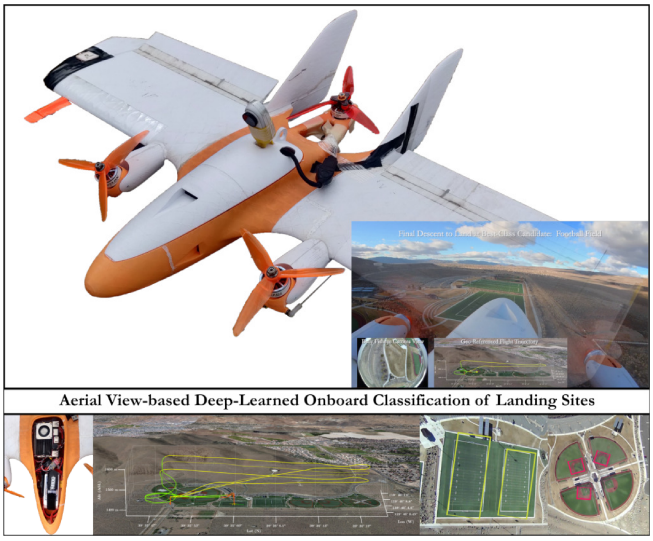

Fig. 1. The MiniHawk-VTOL Tiltrotor aircraft equipped with onboard perception and computing to conduct Deep-Learned aerial view-based detection and classification of landing sites. An indicative field-experiment video sequence is provided as supplementary material.

More specifically, we focus on the critical capacity to perform safe autonomous landing in order to terminate an aerial system mission in the event they are flying above an unmapped location, with possible GPS outage. This work proposes the leveraging of aerial view-based detection and classification of ground landmarks that can act as safe landing sites, towards the purpose of performing landing on-demand without prior knowledge of the underlying terrain structure.

We propose an autonomous landing site discovery pipeline leveraging a Deep Learning-based framework to characterize and identify parts of ground terrain for safe landing spots in an overhead aerial image captured by a Micro Aerial Vehicle (MAV). After detection of potential landing zone candidates, the most promising one for safe landing is selected to terminate the mission. We demonstrate the pipeline by deploying it on an in–house designed Tiltrotor MAV, capable of both Fixed-wing and Vertical Take-Off and Landing (VTOL) in real world settings. Finally, we provide statistical evaluation for the effect of scale and the detection performance at different flying altitudes for the integrated system. Figure 1 depicts the aforementioned platform, the MiniHawk VTOL, carrying the Khadas VIM3 high-level computer that is used in this framework.

* Consider for Best Student Paper Award

This material is based upon work supported by the NSF Award: AWD-01-00002751: RI: Small: Learning Resilient Autonomous Flight. The presented content and ideas are solely those of the authors.

The authors are with the University of Nevada, Reno, 1664 N. Virginia, 89557, Reno, NV, USA prateeka@nevada.unr.edu

The remainder of the paper is organized as follows: Section II presents related work while Section III details the problem setup. Section IV describes system principles and theory, followed by Section V which presents system design and implementation aspects. Section VI demonstrates and discusses field test results. Conclusions are drawn in Section VII.

II. RELATED WORK

Autonomous MAV landing in unknown outdoor as well as indoor environments [21–26] has been widely studied by the robotics community. These works employ advanced control accompanied by estimation and guidance strategy to track and perform precise landing on moving target, by relying on vision based detection of fiducial markers. Authors of [27] propose control algorithm for landing of UAV by utilizing low-cost sensors for precise detection of moving target. The work presented in [28, 29] performs autonomous landing for quadrotor system by relying on onboard sensing and computation. Vlantis et al. [30] tackles the problem of landing a quadrotor on inclined moving platforms by relying on forward–facing camera to detect and track the platform. Furthermore, semi-direct monocular visual odometry (SVO) and simultaneous localization and mapping (SLAM) based approaches such as [31, 32], propose resource–efficient system for real-time mapping of the ground terrain and subsequently detecting landing spot for MAVs by leveraging a 2D height grid map created from the reconstructed terrain. Although, these algorithms are computationally efficient, their deployment on small sized, arm–based resource–constraint companion computer onboard the MAV becomes intractable. Despite the progress, the aforementioned efforts have their own limitations: a) Most existing methods rely on marked based detection of the target landing spots while a few other works utilize additional sensors attached to landing area for detection; b) landing is accomplished within a bounded environment. In this work, we focus on finding suitable landing sites for a Tiltrotor MAV pursuing long distance mission spanning over vast ground terrains following arbitrary path at high altitudes that render vision based mapping techniques ineffective. Our proposed approach relies on the prior knowledge of existence of ground terrain areas such as various sports fields and other level top surfaces that are ubiquitous and present suitable conditions for safe landing. We first detect such suitable candidate landing spots in aerial imagery by leveraging a Deep Learning classification network. Then, the pipeline carries out autonomous navigation to the best landing site and switches to appropriate mode to employ controller to position itself at the geometric center of the landing site. This is followed by initiating a landing command allowing the MAV to descend and finally land at the determined landing spot.

III. PROBLEM SETUP

The problem of landing site discovery in the event of an emergency within the context of this work is considered along the lines of a flying Tiltrotor MAV pursuing an autonomous migratory mission. More specifically, we aim to identify appropriate terrain regions that are safe for landing without creating a comprehensive map of the observable ground terrain and instead leverage a Deep Learning based image classifier to characterize suitable areas in the overheard visual imagery. Considering an autonomous migratory mission that spans over vast arbitrary unmapped and unstructured ground terrain, we distinguish two broad categories of landscape a) uneven elevation regions containing structures that constitute significant collision volume, and b) rigid flat regions with mostly leveled profile such as baseball field, soccer field and structures with flat top surfaces. Within the proposed approach, we aim to characterize and identify the latter class of regions due to their ubiquitous nature and mostly flat surfaces that facilitate safe landing and allow for appropriate takeoff once the batteries are fully charged.

The problems to be addressed are given as follows:

Problem 1 (Vision based landing site detection) It is defined as the problem of bounding-box regression and multi–class terrain landmark classification in overhead aerial images captured via a downward facing camera to find candidate landing sites. The following sub–problem then aims to identify a suitable landing site out of the candidate landing sites.

Problem 2 (Autonomous Tiltrotor Landing) Given a potential landing site and its corresponding georeferenced location ξ^{ref}, the goal is to navigate to the location ξ^{ref} in Fixed-wing configuration and subsequently switch to VTOL configuration to initiate the landing and descent sequences.

A. Our Contribution

The main contribution of our work is summarized as follows:

- We present an integrated pipeline and experimentally verify the applicability of a Deep Learning framework for UAS-based terrain imagery classification.
- We demonstrate how terrain landmarks that lend themselves for safe (emergency) landing are successfully classified in this integrated system design, which can also perform landing on best candidate based on this unlocked level of capacity.
- We provide statistical evaluation for the effect of scale on the accuracy of the detections in MAV UAS imagery acquired across different flying altitudes measured in meters Above Ground Level (AGL).

IV. PROPOSED APPROACH

This section details our contributions towards autonomous landing site discovery for an in–house designed Tiltrotor MAV.

A. Overarching Architecture

For an autonomously navigating aerial vehicle platform with limited power resources, alighting on appropriate terrestrial site is necessary for recharging via solar energy harvesting allowing the vehicle to resume its mission. This is the takeoff point for our pipeline: we consider that the aerial

vehicle is operating in Fixed-wing mode, and we initiate our pipeline as its depleting power resources approach a certain threshold essentially emulating an unexpected event that would require immediate landing. Figure 2 presents a flowchart of this paper's algorithmic components which are subsequently elaborated, illustrating the overarching logic of how these fundamental capabilities contributed through our work can be applied to autonomously detect suitable sites and subsequently land on them.

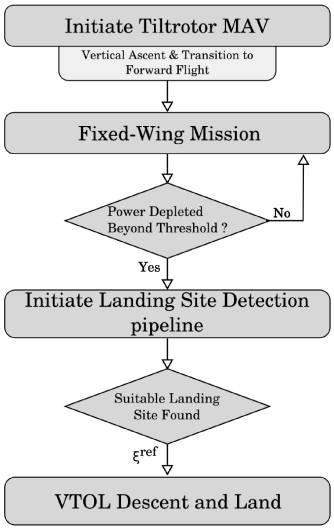

Fig. 2. Flowchart of overarching architecture

B. Vision Based Landing Site Detection

This section presents our implemented pipeline that leverages Deep Learning based object classification network for visual characterization of suitable landing locations in overhead aerial imagery.

We begin by rectifying the fisheye image $\mathbb{I}_f$ by undistoring it with a–priori known intrinsic matrix and distortion parameters to get $\mathbb{I}_U$. It is mentioned here that the camera is well calibrated [33–35] in an offline manner to determine the pinhole camera intrinsic parameters ($[f_u, f_v, p_u, p_v]$) and coefficients of equidistant distortion model ($[k_1, k_2, r_1, r_2]$) that facilitates the undistortion process. We proceed to crop the undistored image $\mathbb{I}_U$ into image sections of dimension 1024×1024, each resulting in four cropped images ($\mathbb{I}_{net_1}, .., \mathbb{I}_{net_4}$) with overlapping regions such that the image dimensions matches the input dimensions of the employed network. An alternate approach to make the image size compatible with the network's input dimension is to resize the image accordingly which however causes a loss in spatial resolution and distorts the visual appearance of the objects in image resulting in failure of detection of even the most obvious object classes. Cropping the original image in smaller overlapping images, on the other hand, ensures that the spatial resolution remains unaffected and high probability of detections due to redundant overlapping pixels. We follow by forward inference of a trained model of object detection network on the cropped images ($\mathbb{I}_{net_1}, .., \mathbb{I}_{net_4}$) to obtain a bounding box and detected object class label. It is noted here that the Deep Learning framework leveraged in this work is capable of classifying a) parts of ground terrain such as baseball–diamond, soccer field, tennis court, basketball court, etc and b) other ground objects such as plane, vehicle, swimming pool, etc. In our approach, we cater to the former category of objects since they are relevant to our use case.

The above mentioned process continuously executes on incoming image frames resulting in detection of candidate landing sites. At the same time, georeferenced location ξ_i^{ref} corresponding to each candidate landing site is identified as well. To qualify as a suitable landing site, we leverage the confidence score associated with each detections provided by the network. During the detection phase we keep a track of all candidate detection and other attributes such as georeferenced location and confidence score. An i^{th} detection with ξ_i^{ref} location with the superior detection score is selected as the final landing site. As soon as the final landing site is identified from the candidate detections, its corresponding georeferenced location ξ^{ref} is forwarded to autonomous landing pipeline which is described in the following section.

Figure 3 depicts an indicative example of the aerial view captured onboard our MAV platform for a sample mission executed in the N. Nevada region. It is highlighted here that our pipeline is able to successfully detect difficult to discern structures such as the storage tank with the sunlight overexposing it, essentially hiding its distinguishing edges and features, and the detection remains consistent while the MAV performs aggressive loiter maneuver which demonstrates the effectiveness of our pipeline.

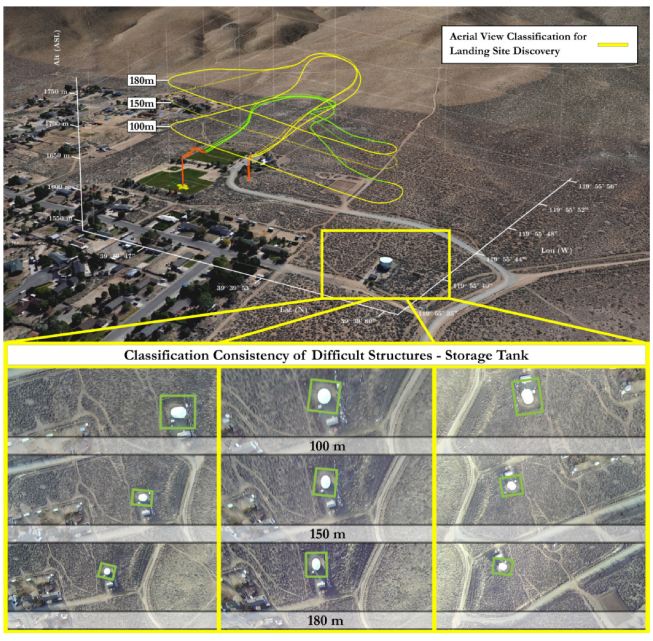

Fig. 3. Top row: Georefernced trajectory of a sample flight path executed in N. Nevada region for aerial view classification of landing site discovery. Bottom rows each representing different flight altitudes, namely $100m$, $150m$ and $180m$ while the columns depict the instances of storage tank detection from different views at each altitudes while performing aggressive loiter circle.

C. Autonomous Tiltrotor landing

For the purposes of this paper we assume that a MAV tasked with an autonomous mission is navigating in Fixed-wing mode through a set of georefenced waypoints. As soon as the power drops below a threshold $\mathcal{T}_w$, the MAV initiates the pipeline described in the previous subsection IV-B to survey the area using a downward facing camera in order to determine mostly planar and safe landing regions. Along the sinuous exploratory path (represented by yellow path in top row of Figure 3 and top row of Figure 5) the algorithmic components described previously characterize several potential landing sites and finally identifies a relevant landing spot with mostly level terrain such as a soccer field. Subsequently, the vehicle navigates to the corresponding georeferenced location ξ_{ref} of the identified landing region and transitions from the Fixed-wing mode to VTOL mode.

Given a body frame of reference $(\mathcal{F}_R)$ rigidly attached to the center of mass of the aircraft in forward–right–down (FRD) configuration, i.e. the x axis aligned along the fuselage pointing towards the nose, the y axis aligned with right wing and the z axis aligned orthogonal to both x and y axis pointing downwards, a PID velocity controller is initiated that operates in plane defined by x and y axis. The controller aims to aligns the geometric center of the detected bounding box $\{x_{bb}, y_{bb}\}$ with the image center $\{x_{img}, y_{img}\}$, both transformed into body frame of reference $(\mathcal{F}_R)$, essentially positioning the vehicle at the center of the landing site at an offset in z axis by its altitude. Immediately after the image center and the bounding box centers are aligned, an autonomous landing sequence is initiated that allows the MAV to descend gradually until it safely lands on the ground.

V. SYSTEM DESIGN

This section details the design of Tiltrotor MAV used in our experiments and the implementation aspect of the proposed approach.

A. Tiltrotor Micro Aerial Vehicle

In this work we use the MiniHawk-VTOL [36, 37] platform, which is a rapidly–prototyped aerial vehicle platform of the operationally versatile [38, 39] Tri-Tiltrotor class [40–44], capable of vertical landing and take off, and Fixed–wing flight for efficient long distance navigation. The platform weighs $1kg$ having $800mm$ wingspan that accommodates solar panels with the ability to cruise at typical speed of $20m/s$. A pair of tiltable BLDC motors actuated by two separate servo motors in the front combined with a fixed rear BLDC provides the utility for VTOL. It is mentioned that the platform we deploy in this work and the proposed algorithmic approach focuses on autonomous vision based landing site discovery rather than the solar harvesting aspect which are discussed in our prior work [37].

In addition to the actuating components, the main body of the vehicle houses several electronic and power related components that are rigidly mounted inside the hollow body including a mRo PixracerPro flight controller and Khadas VIM3 companion computer. The Khadas VIM3 is a powerful

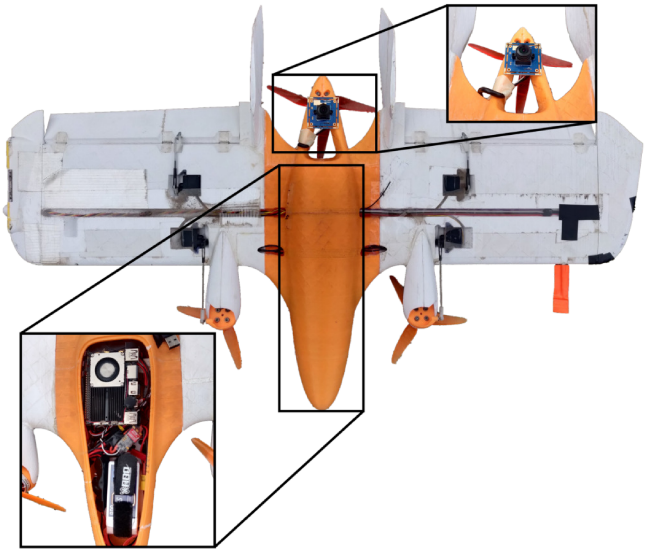

Fig. 4. Bottom view of the MiniHawk VTOL platform. Top right image shows the wide angle 180° field of view camera mounted at the tail end of the fuselage for capturing aerial view imagery. Bottom left image illustrates the top view of the fuselage with the electronic and power components housed in the body of vehicle including Khadas VIM3 companion computer.

small–sized companion computer capable of handling high–level autonomy tasks as well as the ability to run neural network onboard. Packed with a CPU consisting of 4 Cortex-A73 cores paired with 2 Cortex-A53 cores combined with in-built neural processing unit capable of running Deep Neural Networks makes it perfect for our application. Furthermore, it includes a downward facing fisheye lens camera with USB interface and a passive IR-CUT filter that filters out invisible infrared to provide color corrected image in outdoor settings. Figure 4 illustrates the system design; it depicts the image of the bottom view of the vehicle with the downward facing fisheye camera rigidly mounted on the tail that allows us to capture aerial view of terrain during the forward-flight operation.

B. Implementation details

This section describes the implementation specific details of the pipeline that were followed during the experiments.

The MiniHawk VTOL is considered as if operating within the context of an emergency situation, where it has to autonomously decide a safe landing site based in a short time window. Thus, the visual landing site detection pipeline initiates as soon as the vehicle ascends into the air and transitions into forward flight operation, and some surveying is performed prior to detecting a landing site and terminating the mission. We leverage the ArduPilot stack running on the flight controller to both a) perform the initial aerial surveying part of the mission as a waypoint flight trajectory executed in *auto–mode*, as well as b) use the inbuilt velocity controller API to control the horizontal positioning of the vehicle as described in Section IV-C. This feature is available when the vehicle is switched to *guided–mode* and allows the companion computer to guide the MAV to land through the

velocity command interface.

The mission begins with vehicle in auto–mode initiating a VTOL ascent up to a certain altitude, followed by transitioning to forward flight and subsequently navigating as per the waypoint mission. Followingly, the landing site detection pipeline described in Section IV-B is engaged that provides the georeferenced location ξ^{ref} of the landing spot which is set as final navigation waypoint. Immediately after reaching ξ^{ref} the vehicle transitions to VTOL mode and switches to guided–mode allowing the companion computer to position it to the horizontal geometric center of the landing site. At this point, the guided–mode is disabled while the auto–mode is enabled again allowing the companion computer to trigger the landing command.

For the purposes of terrain landmark detection we leverage the YOLOv3 [45] Deep Learning framework, which is capable of discerning multiple instances of object classes in aerial images with densely packed, distributed with large scale variation and arbitrarily oriented objects. In this work, we rely on YOLOv3 network trained on a DOTA dataset [46] containing classes that are relevant to the envisioned problem statement. In order to deploy the trained network model on the NPU on-board Khadas VIM3, it is required to be *quantized*. Quantization refers to the techniques that perform computations and stores the floating point weights to lower bandwidths such as integers. This allows for a more compact representation of the model and the use of high performance vectorized operations on supported hardware without compromising inference time accuracy while significantly reducing computational cost. Mathematically, quantization maps a floating point value $x \in [\alpha, \beta]$ to a b–bit integer $x_b \in [\alpha_b, \beta_b]$ as indicated in the equation 1. To perform model quantization, we leverage npu convert tool [47] which is an SDK that supports models from various frameworks and libraries such as PyTorch, TensorFlow, ONNX, Darknet, etc.

$$x_q = round\left(\frac{1}{s}x + z\right) \tag{1}$$

$$where, \quad s = \frac{\beta - \alpha}{\beta_q - \alpha_q}$$

$$z = round\left(\frac{\beta\alpha_q - \alpha\beta_q}{\beta - \alpha}\right)$$

VI. EXPERIMENTAL STUDY

To experimentally demonstrate the effectiveness of our proposed approach we employed an in–house designed and developed Tiltrotor aerial robot, the MiniHawk-VTOL, armed with the previously presented onboard perception and Deep-Learned classification pipeline that allows it to detect ground landmarks of significance.

The presented test were conducted in the region of N. Nevada at three different cruising altitudes of $100m$, $150m$ and $180m$ to demonstrate the efficacy of our approach. Figure 5 depicts the corresponding georeferenced flight trajectory within the context of such a landing-on-request mission. Key instances of candidate landing regions are annotated on the figure, as detected by the Deep Learning framework, and associated to the corresponding sections of the vehicle's flight path. The mission is designed to emulate a search for candidate landing sites due to an unexpected event. Therefore, the flight sequence comprises of three phases, namely i) VTOL altitude ascent to $60m$ above ground level, ii) Fixed-wing climb to cruising altitude and surveying for landing site detection, initiated due to the detection of an envisioned emergency event which requires to early-terminate the mission, and finally iii) a VTOL descent and landing at an identified suitable site discovered along this search.

The specific environment terrain for the mission comprises of various sport fields and other cluttered and densely packed objects. The mission process follows the outline presented in Section IV-A; the vehicle ascends up to an altitude of $60m$, indicated by the orange path in Figure 5 followed by transitioning into forward flight configuration to climb to the cruising altitude and surveying in Fixed–wing mode. At the same time the Deep Leaning detection pipeline is engaged which begins to characterize and detect various candidate landing sites. The bottom row of Figure 5 depicts instances of baseball diamonds and soccer fields detected along the flight path. These are illustrated with association to the corresponding sections of the aerial robot's trajectory.

Following this, a suitable landing spot out of the ones detected earlier is identified along with the georeferenced location. In this particular mission the soccer field indicated in the Figure 5-c) is selected as an appropriate landing site, based on its class and superior detection score. Subsequently, the MAV navigates toward the corresponding location and transitions into VTOL operation mode allowing it to land at the selected landing spot.

It is noted that an indicative video sequence of the $180m$ AGL flight is provided as supplementary material with this paper.

A. Statistical Evaluation

In this sub–section we discuss the performance of the Deep Learning network and the effect of scale for the classification of various landmarks in continuous sequence of aerial imagery. Table I provides precision, recall and detection performance metrics for the experimental tests conducted at two different sites in the N. Nevada region across different flying altitudes measured above ground level (as indicated in Figures 3 and 5). It is mentioned here that the statistical results for Silver Knolls Park (SKP) flight sequence (in I) are provided with respect to a single class (storage tank) of a difficult to detect structure during aggressive loitering of MAV while the Golden Eagle Park (GEP) flight sequence contains results for multi–class classification.

The overall detection accuracy follows an inverse relation with the scale of the landmarks (proportional relation with the AGL flying altitude) while precision and recall follow the same trend for GEP flight sequences. However, for SKP flight sequence images containing relatively small structure such as storage tank, the detection accuracy, precision and recall, all decrease with the decrease in scale. This is

TABLE I

PERFORMANCE OF THE DEEP LEARNING CLASSIFICATION NETWORK COMPARED FOR THREE IMAGE SCALES CAPTURED ACROSS DIFFERENT AGL FLYING ALTITUDES AT TWO SITES IN THE N. NEVADA REGION.

Flight sequence	True Positive (TP)	False Positive (FP)	False Negative (FN)	Detection Performance	Precision	Recall
Golden Eagle Park - 100m alt	613	159	852	37.74%	0.794	0.418
Golden Eagle Park - 150m alt	783	67	1282	36.72%	0.921	0.379
Golden Eagle Park - 180m alt	857	77	1098	42.17%	0.917	0.438
Silver Knolls Park - 100m alt	37	4	34	49.33%	0.902	0.521
Silver Knolls Park - 150m alt	32	3	69	30.76%	0.914	0.316
Silver Knolls Park - 180m alt	51	8	103	31.48%	0.864	0.331

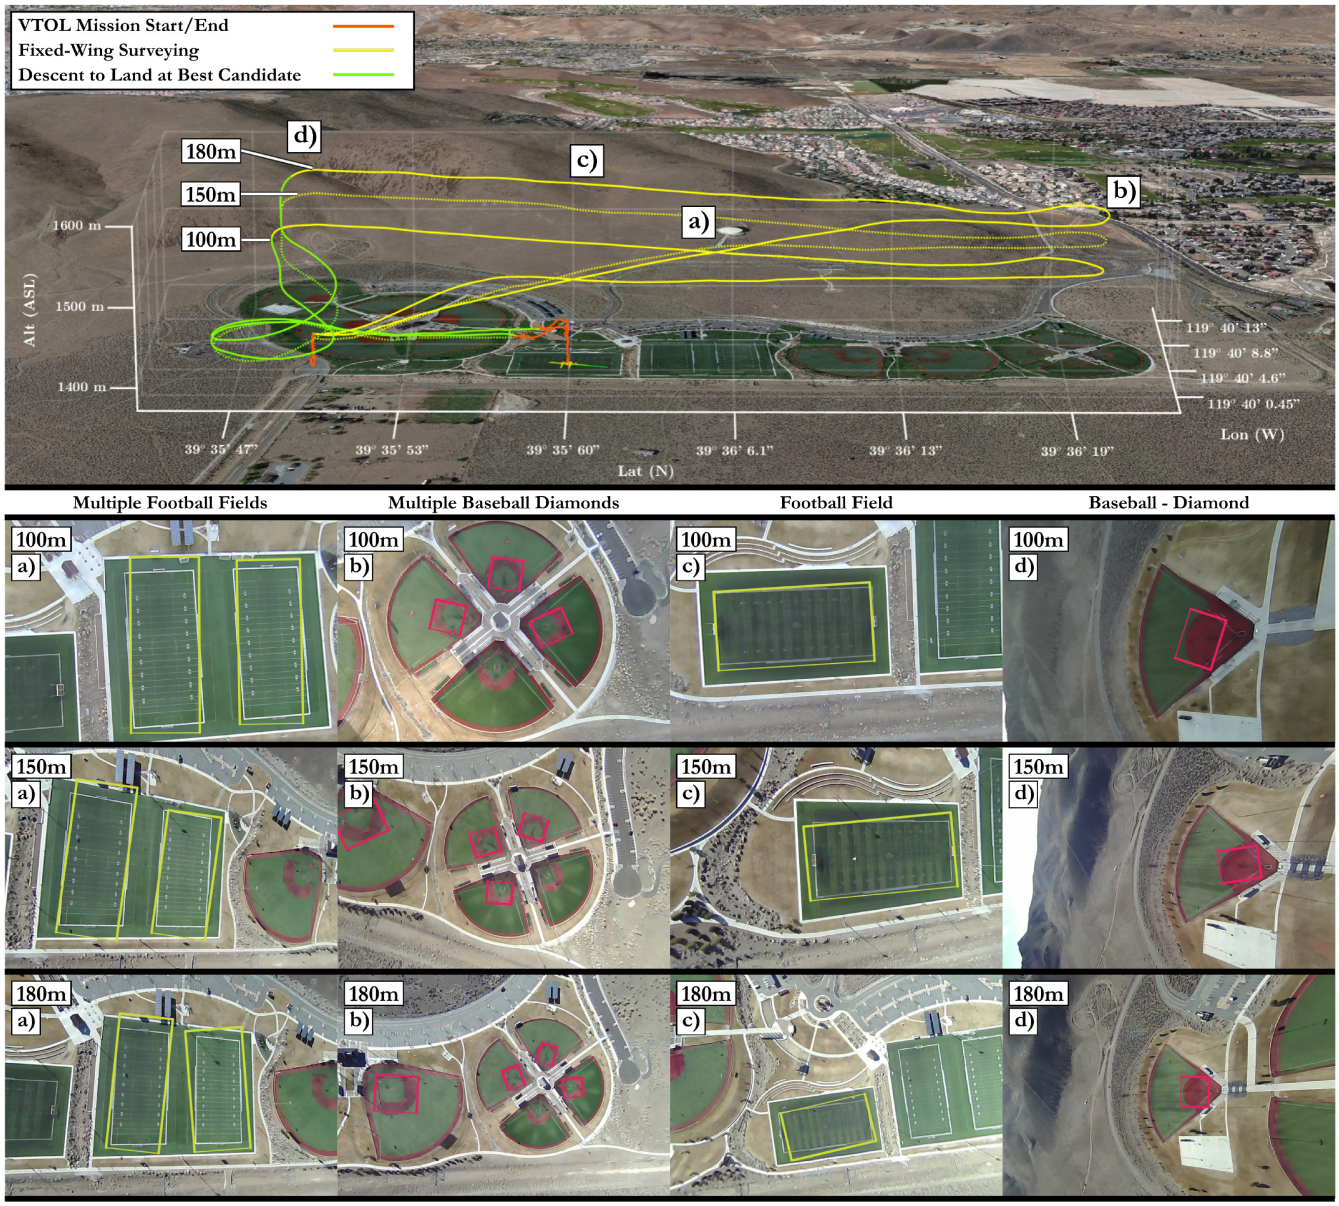

Fig. 5. Top row: Mission phases illustrated as georeferenced flight paths at three distinct AGL altitudes, namely $100m$, $150m$ and $180m$, overlaid on a perspective view of a N. Nevada region. Orange-colored: Mission start with VTOL ascent, mission end with VTOL landing. Yellow-colored: Fixed-wing flight with simultaneous aerial classification of landing sites. Green-colored: Fixed-wing final approach of best candidate for landing. The three bottom rows correspond to the Fixed–wing AGL flight altitude in ascending order of $100m$, $150m$ and $180m$, while the columns represents key detection instances; a) an instance of multiple football fields, b) an instance of multiple baseball diamonds, c) soccer field selected as best candidate for landing site, d) single baseball diamond detection during sharp loiter turns. An indicative video sequence of the experiment is provided as supplementary material with this paper.

an expected behaviour for such small structures since the decrease in scale results in the loss of spatial resolution and essentially lower classification confidence score for the storage tank class. Indicatively, altitude of $180m$, despite

having slightly low accuracy as compared to $150m$ altitude flight, is excellently tailored for a UAS terrain landmark classification purposes providing a balanced tradeoff between precision and recall for all the classes of landing sites.

VII. CONCLUSIONS

In this work, a vision-based pipeline to discover landing sites based on a deep-learned classification policy, and by relying only on aerial imagery captured onboard an autonomous aerial system was proposed and demonstrated. The proposed policy contributes an effective pipeline for characterizing potential landing regions and essentially identifying the best candidate for a landing of the vehicle while autonomously landing at the corresponding georeferenced location in a VTOL manner. The proposed approach was evaluated and tested in real-world settings within the context of multiple field tests conducted in the Northern Nevada region demonstrating its consistency and effectiveness.

REFERENCES

[1] N. Michael, S. Shen, K. Mohta, Y. Mulgaonkar, V. Kumar, K. Nagatani, Y. Okada, S. Kiribayashi, K. Otake, K. Yoshida *et al.*, "Collaborative mapping of an earthquake-damaged building via ground and aerial robots," *Journal of Field Robotics*, vol. 29, no. 5, pp. 832–841, 2012.

[2] M. Tranzatto, F. Mascarich, L. Bernreiter, C. Godinho, M. Camurri, S. M. K. Khattak, T. Dang, V. Reijgwart, J. Loeje, D. Wisth, S. Zimmermann, H. Nguyen, M. Fehr, L. Solanka, R. Buchanan, M. Bjelonic, N. Khedekar, M. Valceschini, F. Jenelten, M. Dharmadhikari, T. Homberger, P. De Petris, L. Wellhausen, M. Kulkarni, T. Miki, S. Hirsch, M. Montenegro, C. Papachristos, F. Tresoldi, J. Carius, G. Valsecchi, J. Lee, K. Meyer, X. Wu, J. Nieto, A. Smith, M. Hutter, R. Siegwart, M. Mueller, M. Fallon, and K. Alexis, "Cerberus: Autonomous legged and aerial robotic exploration in the tunnel and urban circuits of the darpa subterranean challenge," *Journal of Field Robotics*, 2021.

[3] T. Tomic, K. Schmid, P. Lutz, A. Domel, M. Kassecker, E. Mair, I. L. Grixa, F. Ruess, M. Suppa, and D. Burschka, "Toward a fully autonomous uav: Research platform for indoor and outdoor urban search and rescue," *IEEE robotics & automation magazine*, vol. 19, no. 3, pp. 46–56, 2012.

[4] P. Arora and C. Papachristos, "Mobile manipulation–based deployment of micro aerial robot scouts through constricted aperture-like ingress points," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 6716–6723.

[5] M. Burri, J. Nikolic, C. Hürzeler, G. Caprari, and R. Siegwart, "Aerial service robots for visual inspection of thermal power plant boiler systems," in *2012 2nd international conference on applied robotics for the power industry (CARPI)*. IEEE, 2012, pp. 70–75.

[6] A. Bircher, K. Alexis, M. Burri, P. Oettershagen, S. Omari, T. Mantel and R. Siegwart, "Structural inspection path planning via iterative viewpoint resampling with application to aerial robotics," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 6423–6430.

[7] C. Papachristos, K. Alexis, L. R. G. Carrillo, and A. Tzes, "Distributed infrastructure inspection path planning for aerial robotics subject to time constraints," in *2016 international conference on unmanned aircraft systems (ICUAS)*. IEEE, 2016, pp. 406–412.

[8] F. Mascarich, T. Wilson, C. Papachristos, and K. Alexis, "Radiation source localization in gps-denied environments using aerial robots," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6537–6544.

[9] S. Khattak, C. Papachristos, and K. Alexis, "Visual-thermal landmarks and inertial fusion for navigation in degraded visual environments," in *2019 IEEE Aerospace Conference*. IEEE, 2019, pp. 1–9.

[10] C. Papachristos and K. Alexis, "Augmented reality-enhanced structural inspection using aerial robots," in *2016 IEEE international symposium on intelligent control (ISIC)*. IEEE, 2016, pp. 1–6.

[11] S. Khattak, C. Papachristos, and K. Alexis, "Keyframe-based direct thermal–inertial odometry," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3563–3569.

[12] S. Khattak, C. Papachristos, and K. Alexis, "Keyframe-based thermal–inertial odometry," *Journal of Field Robotics*, vol. 37, no. 4, pp. 552–579, 2020.

[13] S. Khattak, F. Mascarich, T. Dang, C. Papachristos, and K. Alexis, "Robust thermal-inertial localization for aerial robots: A case for direct methods," in *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2019, pp. 1061–1068.

[14] T. Dang, F. Mascarich, S. Khattak, H. Nguyen, N. Khedekar, C. Papachristos, and K. Alexis, "Field-hardened robotic autonomy for subterranean exploration," *Field and Service Robotics (FSR)*, 2019.

[15] C. Papachristos, S. Khattak, and K. Alexis, "Uncertainty-aware receding horizon exploration and mapping using aerial robots," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 4568–4575.

[16] C. Papachristos, M. Kamel, M. Popović, S. Khattak, A. Bircher, H. Oleynikova, T. Dang, F. Mascarich, K. Alexis, and R. Siegwart, "Autonomous exploration and inspection path planning for aerial robots using the robot operating system," in *Robot Operating System (ROS)*. Springer, Cham, 2019, pp. 67–111.

[17] C. Papachristos, F. Mascarich, S. Khattak, T. Dang, and K. Alexis, "Localization uncertainty-aware autonomous exploration and mapping with aerial robots using receding horizon path-planning," *Autonomous Robots*, vol. 43, no. 8, pp. 2131–2161, 2019.

[18] W. Johnson, S. Withrow-Maser, L. Young, C. Malpica, W. Koning, K. WJF, M. Fehler, A. Tuano, A. Chan, A. Datta *et al.*, *Mars Science Helicopter Conceptual Design*. National Aeronautics and Space Administration, Ames Research Center, 2020.

[19] R. D. Lorenz, E. P. Turtle, J. W. Barnes, M. G. Trainer, D. S. Adams, K. E. Hibbard, C. Z. Sheldon, K. Zacny, P. N. Peplowski, D. J. Lawrence *et al.*, "Dragonfly: A rotorcraft lander concept for scientific exploration at titan," *Johns Hopkins APL Technical Digest*, vol. 34, no. 3, p. 14, 2018.

[20] S. J. Carlson and C. Papachristos, "Solar energy harvesting for a land–to–recharge tiltrotor micro aerial vehicle," in *2022 IEEE Aerospace Conference (AEROCONF)*, 2022.

[21] S. Lee, T. Shim, S. Kim, J. Park, K. Hong, and H. Bang, "Vision-based autonomous landing of a multi-copter unmanned aerial vehicle using reinforcement learning," in *2018 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2018, pp. 108–114.

[22] H. Lee, S. Jung, and D. H. Shim, "Vision-based uav landing on the moving vehicle," in *2016 International conference on unmanned aircraft systems (ICUAS)*. IEEE, 2016, pp. 1–7.

[23] J. Kim, Y. Jung, D. Lee, and D. H. Shim, "Outdoor autonomous landing on a moving platform for quadrotors using an omnidirectional camera," in *2014 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2014, pp. 1243–1252.

[24] Y. Jung, D. Lee, and H. Bang, "Close-range vision navigation and guidance for rotary uav autonomous landing," in *2015 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2015, pp. 342–347.

[25] O. Araar, N. Aouf, and I. Vitanov, "Vision based autonomous landing of multirotor uav on moving platform," *Journal of Intelligent & Robotic Systems*, vol. 85, no. 2, pp. 369–384, 2017.

[26] R. Polvara, S. Sharma, J. Wan, A. Manning, and R. Sutton, "Vision-based autonomous landing of a quadrotor on the perturbed deck of an unmanned surface vehicle," *drones*, vol. 2, no. 2, p. 15, 2018.

[27] Y. Jung, S. Cho, and D. H. Shim, "A trajectory-tracking controller design using l 1 adaptive control for multi-rotor uavs," in *2015 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2015, pp. 132–138.

[28] D. Falanga, A. Zanchettin, A. Simovic, J. Delmerico, and D. Scaramuzza, "Vision-based autonomous quadrotor landing on a moving platform," in *2017 IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*. IEEE, 2017, pp. 200–207.

[29] J. Kim, Y. Jung, D. Lee, and D. H. Shim, "Landing control on a mobile platform for multi-copters using an omnidirectional image sensor," *Journal of Intelligent & Robotic Systems*, vol. 84, no. 1, pp. 529–541, 2016.

[30] P. Vlantis, P. Marantos, C. P. Bechlioulis, and K. J. Kyriakopoulos, "Quadrotor landing on an inclined platform of a moving ground vehicle," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 2202–2207.

[31] C. Forster, M. Faessler, F. Fontana, M. Werlberger, and D. Scaramuzza, "Continuous on-board monocular-vision-based elevation mapping applied to autonomous landing of micro aerial vehicles," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 111–118.

[32] T. Yang, P. Li, H. Zhang, J. Li, and Z. Li, "Monocular vision slam-based uav autonomous landing in emergencies and unknown environments," *Electronics*, vol. 7, no. 5, p. 73, 2018.

[33] L. Oth, P. Furgale, L. Kneip, and R. Siegwart, "Rolling shutter camera calibration," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 1360–1367.

[34] J. Maye, P. Furgale, and R. Siegwart, "Self-supervised calibration for robotic systems," in *2013 IEEE Intelligent Vehicles Symposium (IV)*, 2013, pp. 473–480.

[35] P. Furgale, T. D. Barfoot, and G. Sibley, "Continuous-time batch estimation using temporal basis functions," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 2088–2095.

[36] S. J. Carlson and C. Papachristos, "The minihawk-vtol: Design, modeling, and experiments of a rapidly-prototyped tiltrotor uav," in *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2021, pp. 777–786.

[37] S. J. Carlson and C. Papachristos, "Migratory behaviors, design principles, and experiments of a vtol uav for long-term autonomy," *ICRA 2021 Aerial Robotics Workshop on "Resilient and Long-Term Autonomy for Aerial Robotic Systems"*, 2021.

[38] C. Papachristos, K. Alexis, and A. Tzes, "Technical activities execution with a tiltrotor uas employing explicit model predictive control," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 11 036–11 042, 2014.

[39] S. J. Carlson, P. Arora, and C. Papachristos, "A multi–vtol modular aspect ratio reconfigurable aerial robot," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, 2022.

[40] C. Papachristos, K. Alexis, and A. Tzes, "Design and experimental attitude control of an unmanned tilt-rotor aerial vehicle," in *2011 15th International Conference on Advanced Robotics (ICAR)*. IEEE, 2011, pp. 465–470.

[41] C. Papachristos and A. Tzes, "Modeling and control simulation of an unmanned tilt tri-rotor aerial vehicle," in *2012 IEEE International Conference on Industrial Technology*. IEEE, 2012, pp. 840–845.

[42] C. Papachristos, K. Alexis, and A. Tzes, "Towards a high-end unmanned tri-tiltrotor: Design, modeling and hover control," in *2012 20th Mediterranean Conference on Control & Automation (MED)*. IEEE, 2012, pp. 1579–1584.

[43] C. Papachristos, K. Alexis, and A. Tzes, "Model predictive hovering-translation control of an unmanned tri-tiltrotor," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 5425–5432.

[44] C. Papachristos, K. Alexis, and A. Tzes, "Dual–authority thrust–vectoring of a tri–tiltrotor employing model predictive control," *Journal of intelligent & robotic systems*, vol. 81, no. 3-4, pp. 471–504, 2016.

[45] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.

[46] G.-S. Xia, X. Bai, J. Ding, Z. Zhu, S. Belongie, J. Luo, M. Datcu, M. Pelillo, and L. Zhang, "Dota: A large-scale dataset for object detection in aerial images," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3974–3983.

[47] K. T. Frank Wang, "Npu convert tool sdk," https://github.com/khadas/aml_npu_sdk, 2019.