AirDrop: Towards Collaborative, Multi-Resolution Air–Ground Teaming for Terrain-Aware Navigation

Kasthuri Jayarajah†, Sean Gart‡, Aryya Gangopadhyay†
† University of Maryland, Baltimore County, ‡DEVCOM Army Research Lab

ABSTRACT

Driven by advances in deep neural network models that fuse multi-modal input such as RGB and depth representations to accurately understand the semantics of the environments (e.g., objects of different classes, obstacles, etc.), ground robots have gone through dramatic improvements in navigating unknown environments. Relying on their singular, limited perspective, however, can lead to suboptimal paths that are wasteful and quickly drain out their batteries, especially in the case of long-horizon navigation. We consider a special class of ground robots, that are air-deployed, and pose the central question: can we leverage aerial perspectives of differing resolutions and fields of view from air–to–ground robots to achieve superior terrain-aware navigation? We posit that a key enabler of this direction of research is collaboration between such robots, to collectively update their route plans, leveraging advances in long-range communication and on-board computing. Whilst each robot can capture a sequence of high resolution images during their descent, intelligent, lightweight pre-processing on-board can dramatically reduce the size of the data that needs to be shared among its peers over severely bandwidth-limited long range communication channels (e.g., over sub gigahertz frequencies). In this paper, we discuss use cases and key technical challenges that must be resolved to realize our vision for **collaborative, multi-resolution terrain-awareness for air–to–ground robots**.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**;

KEYWORDS

Edge Intelligence, Autonomous Systems

ACM Reference Format:
Kasthuri Jayarajah†, Sean Gart‡, Aryya Gangopadhyay† . 2023. *AirDrop*: Towards Collaborative, Multi-Resolution Air–Ground Teaming for Terrain-Aware Navigation. In *The 24th International Workshop on Mobile Computing Systems and Applications (HotMobile '23), February 22–23, 2023, Newport Beach, CA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3572864.3580335

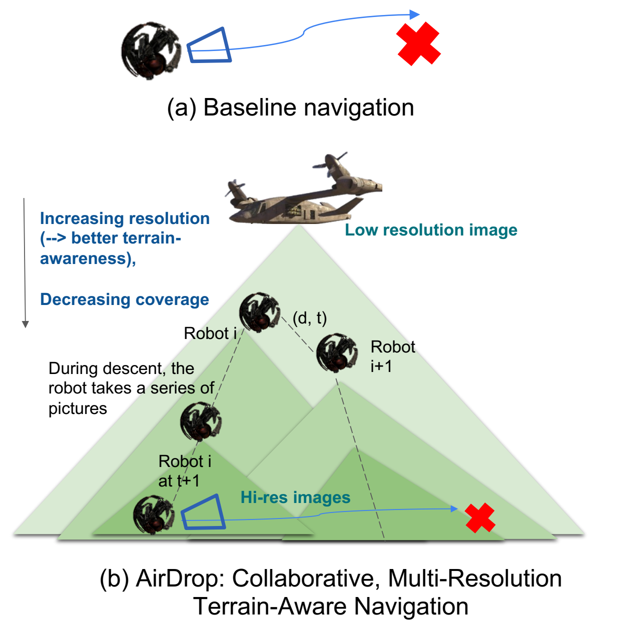

Figure 1: Contrasting (a) traditional ground navigation based on local sensing against (b) *AirDrop*-based navigation with local + global terrain awareness

1 INTRODUCTION

Recent advances in algorithms for autonomous navigation [9] together with mobile computers that can now run state-of-the-art perception models, have led to ground robots as well as aerial robots being deployed in a variety of practical situations (e.g., in manufacturing environments [4], terrain-aware search and rescue [16], etc.). Very recently, efforts towards terrain-aware path planning based on techniques such as reinforcement learning [15, 36, 39] have gained significant traction where robots have better situational awareness leading to better path choices. However, singular perspectives of such ground robots, may lead to local decisions that may turn out to be expensive, especially in scenarios where the robots need to navigate over large areas with differing terrains. To provide a global perspective of the terrain, recent works have also considered the combination of air–ground perspectives where one or more aerial and ground robot partners work together [5, 13]. However, these works are limited by the low-resolution nature of the high altitude aerial perspectives and system evaluations do not consider wide area navigation. Our preliminary evaluations over benchmark datasets show that with a decrease in resolution (from 200 to 100 pixels, for e.g.), the object detection accuracy using state-of-the-art deep neural networks (DNNs) dramatically drop by up to 15%.

HotMobile '23, February 22–23, 2023, Newport Beach, CA, USA

ACM ISBN 979-8-4007-0017-0/23/02…$15.00
https://doi.org/10.1145/3572864.3580335

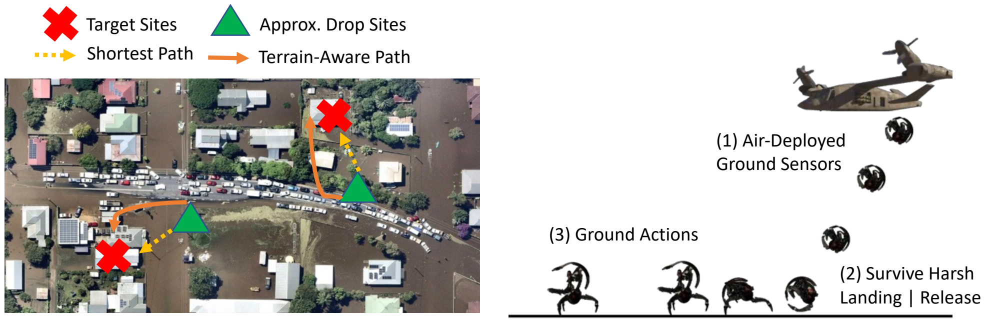

Figure 2: Illustrative Scenarios of *AirDrop*. Left–Air-delivering humanitarian aid during disaster response, Right–Air-deploying ground sensors in contested environments.

In this position paper, we present a novel paradigm for **air–ground teaming**, illustrated in Figure 1; where multiple, collaborative robots that are *air-delivered*, communicate over long-range radio, to share among themselves *light-weight aerial digests* of increasing resolution (of narrower regions) ascertained during their descent. These digests are then fused with low resolution, albeit wide area, perspectives to compute more accurate terrain cost maps used in generating ground navigation paths for the cooperating robots. After reaching the ground, these robots then combine their locally sensed ground situation together with their collaborative knowledge of the global terrain for making better navigation choices. This line of enquiry requires tackling multiple design and system challenges for making collaborative aerial perception feasible on severely resource-constrained devices (e.g., communication in the sub gigahertz frequencies over multiple kilometers limited to ≤ 10 kBits/second).

Our goal for this paper is not to discuss specifics of system design, but to motivate the novel setting of collaborative, distributed aerial perception for superior ground navigation. In the remainder of the paper, we describe scenarios for which *AirDrop* is designed for, goals guiding the architecture of *AirDrop* and discuss key technical challenges that must be addressed in realizing *AirDrop*.

2 MOTIVATING SCENARIOS

We discuss two future scenarios, illustrated in Figure 2, that *AirDrop* supports.

Air-deploying ground sensors in contested environments: In contested environments, ground sensors used for zonal reconnaissance are typically deployed by human soldiers which leaves them exposed to adversarial forces for prolonged periods. Since recently, air-delivering sensors using unmanned aerial vehicles has gained traction – however, in settings such as these – this poses two main challenges: (i) the electronics onboard the drones, especially during active flight, emit radio frequency signatures [30] that can lead to them being easily detected, and (ii) the maximum payload that can be carried is limited (e.g., maximum of 2.5 kg supported by the Mikrokopter [25]). In order to remain discreet, an alternate mode of air–delivering is to *drop* the sensors, or robotic agents, from air where external structures surrounding the robot can absorb the forces from the impact with the ground and protect the robot within. The descent then is governed by the laws of free fall. Once reaching ground, mechanisms on the robot allow it to *escape* and continue with its ground mission (e.g., clearing vegetation at specific locations, station itself at specific locations for surveillance, etc.). Based on the aerial perspectives ascertained during descent, global reasoning around the terrain characteristics in the region (e.g., grassy patches and water bodies that are costly to navigate) together with the locally sensed surrounding (e.g., immediate obstacles) guide the ground robots in navigating towards their target sites.

Last-mile humanitarian aid delivery during disaster response: Air-delivering humanitarian aids (such as dry rations, vaccines, etc.) has been a common last-mile distribution mechanism as part of disaster response operations. However, delivering from higher altitudes result in the aid being sometimes dropped off in inaccessible sites. On the contrary, air-dropping cheap yet effective ground robots carrying such payloads, electronically designed to maneuver their descent (e.g., acceleration-triggered tilting of the robot body to move left or right), can overcome such challenges. From an approximate drop site, as illustrated in Figure 2, the robots can then navigate over less costly routes to reach target site(s) as compared to expensive, or sometimes inaccessible, shortest paths.

In both these scenarios, we make the fair assumption that the environments in which the sensors/robots are being deployed in are (i) highly dynamic (e.g., hostile, post-disaster, etc.) and (ii) require that the robots navigate towards targets that are often mobile (e.g., people). This serves as the basis of the collaborative approach we articulate in this paper where we envision multiples of custom-deployed robots cooperate for high fidelity imaging. We also point out to the astute reader, in situations where these is no requirement to be discreet, a single UAV flying at various altitudes can be an alternative for spatially distributed, multi-resolution sensing of the terrain. The downstream technical challenges as we describe later in Section 3.3, nevertheless, remain relevant.

3 DESIGNING *AIRDROP*

As we introduced in Section 1, state-of-the-art models in ground navigation rely on their local perspectives potentially leading to suboptimal paths on wide-area navigation tasks. In this section, we identify salient features of a collaborative, distributed, aerial perception paradigm, discuss a preliminary architecture for *AirDrop*, and discuss key technical challenges.

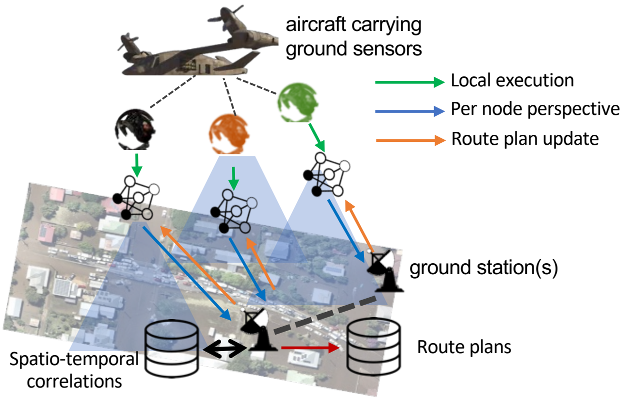

Figure 3: End-to-end, collaborative terrain-awareness architecture

3.1 Opportunities for Collaborative, Terrain-Awareness

We posit that a collaborative, air–to–ground approach provides significant benefits over a non-collaborative, local approach:

Spatially distributed, multi-resolution perspectives: While wide area perspectives can provide an overview of the vast terrain in a few images, as we see in section 4.1, the lower resolution images suffer from significant deterioration in performance for vision/perception models. On the other hand, aerial perspectives that can be taken at lower altitudes can produce higher resolution images but at a loss of coverage. Hence, we posit that a set of spatially distributed air-born sensors/robots can provide both the desired coverage and resolution.

Accurate terrain awareness: As a result of the above higher quality images, and better performance of vision models, reasoning over terrain for wide areas can become more accurate, for the set of robots. On the contrary, robots that rely only on local perspectives of their own and operate in isolation, are severely disadvantaged.

3.2 *AirDrop*: Design Principles

In Figure 3, we illustrate a preliminary architecture for *AirDrop*, and articulate the following system components.

Central server: Initially, we assume that at least a single, wide area coverage aerial image of the area (e.g., satellite imagery or taken by the aircraft delivering the ground robots) to be surveyed exists. We envision that one (or more) ground stations that acts as the central server runs the baseline path planning for the collective of robots on mission. It then receives aerial perspectives of increasing resolution from the individual robots during their descent and by establishing spatio-temporal correlations amongst them, it stitches the pieces of subregions covered by the images to create a super-resolution image that it uses to run algorithms such as image segmentation for reasoning over the terrain underneath. Based on the intermittent updates from the spatially distributed robots, it then revises the set of planned paths and updates over long range communication. A key design goal here is that the more expensive, compute-intensive operations are offloaded to a more powerful compute node while having the edge node (robots in this case) carry out only lighter tasks. In the event that ground stations do not exist, we assume that the aircraft is large enough to match the computational and power needs to receive image updates from the individual robots and update the path plans.

Air–to–ground robots: During descent, the edge nodes intermittently take aerial perspective shots, and upload them back to the *server*. These nodes are also equipped with onboard compute capabilities where they can run full/partial neural network models. However, a key goal here is to minimize the resource consumption during the descent to maximize resources availability and operational time during the actual ground navigation. After reaching the ground, together with the updated paths and real-time situation awareness on the ground (i.e., the local perspective), the robots then navigate towards predefined targets.

Alternative designs: Satellite imagery has demonstrated to be useful in providing aerial perspectives for a variety of applications (e.g., monitoring climate change [27], wildfire flareups [20] and flood forecasting [38]). However, recent works have shown that (i) detecting small objects such as vehicles or people, especially when not present in clear and vast fields, is difficult [14], and (ii) such high-quality satellite imagery with short re-visit times are not easily accessible to disaster-prone, developing regions [28] to enable quick response. *AirDrop's* design towards acquiring high-resolution, real-time aerial perspectives is thus motivated by the scenarios we consider – where targets such as people are mobile and the ground situation is highly dynamic (e.g., post-disaster, hostile environments, etc.).

Metrics of interest: A key performance measure for *AirDrop*, is the improvement in navigation efficacy – e.g., total trip time and energy consumption for long-horizon navigation tasks, as compared to baseline models that only leverage (i) ground perspectives and (ii) a single aerial, low resolution perspective in combination with ground. Furthermore, the overheads in enabling multi-resolution collaboration need to be accounted for; for e.g., the additional energy consumption from serial image acquisition during the descent, any onboard computation, and due to any algorithmic changes to how they read the terrain cost maps together with the ground perspectives sensed.

3.3 Technical Challenges

In this section, we highlight key technical challenges that must be addressed to fully exploit collaborative, multi-resolution scanning.

1. Optimally deploying ground robots: To cover a given area with a minimum guaranteed accuracy, *AirDrop* should devise an optimal strategy for air-deploying the robots. The spatio-temporal scheduling of the individual robot drops should account for several factors including the number of robots, the sensing fidelity vs. field of view vs. altitude characteristics, and fall characteristics (e.g., horizontal departure from an estimated drop site due to prevailing winds). While existing optimization frameworks [6, 10] have investigated optimal placement of static sensors in wireless sensor networks, devising a strategy for coverage coupled together with the physical characteristics of the descent is an interesting problem to pursue.

2. Tradeoffs in multi-resolution semantic segmentation: Furthermore, from the series of multi-resolution images collaboratively produced, *AirDrop stitches* a higher quality, fused image over which it runs instance segmentation. Labels from the segmentation step are then used to produce terrain cost maps that guide path planning of the collaborating robots. While a large number of images can lead to more accurate segmentation inferences, the compute-intensive nature of the semantic segmentation process, as fully-convolutional layers are executed (as part of DNNs) to output per-pixel level labels, presents an interesting trade-off between accuracy and execution efficiency.

3. Lightweight state sharing: In the settings that we consider in this paper, the robots are deployed over a wide area requiring long range communication for any data or state sharing between them. Measurement studies [40] show that even at the highest transmission power (e.g., 27 dBm), a nominal increase of the bitrate to 9.6 kbit/s or more can reduce the range of communication in the sub gigahertz frequencies from 8 to 6 kilometers. On the other hand, typical images of resolution 1280 × 960 can easily take up to 4.7 Mb in size[1] translating to ≈8 minutes to transfer over such channels, possibly with high packet losses. Prior works in 2D image processing have looked at prioritization of image regions based on criticality [22, 23] and overlap [34]. In these works, the content of the image changes lending selection of subregions feasible. On the contrary, in our setting, the content of the image is expected to remain the same, however, revealing more salient attributes that might change the learned semantics of certain subregions as the resolution increases. This poses an interesting challenge where potential approaches involving partial or local computation onboard the robot, akin to generating perspective digests from early layers of image processing pipelines [19], may reduce the size of the data that is shared drastically.

4. Surviving harsh landings and release mechanisms: As we discuss in Section 2, for contested environments where deployments need to be discreet (fast and minimal radio frequency signatures to minimize probability of detection by adversarial forces), an additional design challenge is in creating mechanisms for the robots to survive the landings on ground (after being air-deployed). While crumple zone-like designs [26] may be effective, a related challenge is in designing appropriate mechanisms for the robot to then release itself from the contraptions that help them survive the landing.

4 PRELIMINARY EVALUATION

In this section, we first provide preliminary quantitative evidence of how the resolution of an image impacts the downstream accuracy of image processing tasks, taking object detection as an exemplar. We then establish the worst case speeds air-dropped sensors would undergo during their descent which can challenge the stability of the images captured.

4.1 Resolution Vs. Accuracy Tradeoff

Here, we report on our preliminary experiments on quantifying the degradation of accuracy in analysing images with increasing altitude. Due to the lack of labelled vision datasets at varying altitudes, we create a synthetic, multi-resolution dataset from the existing single resolution benchmark dataset, PETS2009 [12], by systematically reducing the resolution to various extents (to proxy altitude change).

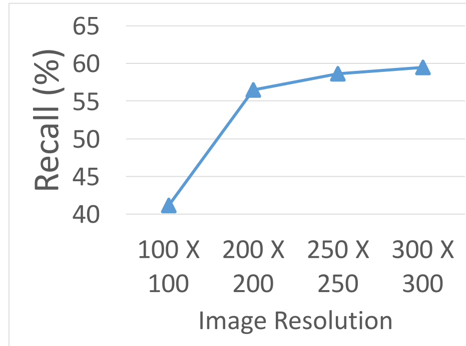

Figure 4: Image resolution vs. object detection recall on the PETS2009 dataset.

In Figure 4, we plot the recall of running a widely used, lightweight object detection DNN, SSD [24], on the dataset (on the $y-$ axis) against the resolution of the image on the $x-$ axis. We keep the *coverage* of the images the same across the different resolution cases. We see that while the model is resilient to resolutions close to its input size (i.e., 300 × 300, the model suffers a dramatic reduction in performance (i.e., from 56% to 40%) when the input resolution drops from 200 to 100 pixels.

4.2 Worst Case Speeds During Descent

One of the key design ideas behind the proposed system is to ascertain images, in increasing resolution, during the descent. However, the speed of the descent, and the resulting stability, will impact the legibility of the images. We consider a legged hexapod suitable for air-dropping which weights approximately 6 Lbs. (≈2.2 kg), contains a Raspberry Pi 4 onboard for computation, and an operation time of roughly 40 minutes (see Figure 5)[2]. Furthermore, the hexapod can also support a payload of up to 1 kg which makes it appropriate for transporting emergency relief (e.g., dry rations and medicines) in the post-disaster response scenario we describe earlier. Together with the Raspberry Pi 4 board, each unit costs ≤500 USD. We calculate the terminal velocity during the descent assuming that the robot undergoes free fall – i.e., devoid of any mechanisms for a controlled flight – as the worst case speed using the formula $V_t = \sqrt{\frac{2mg}{\rho A C_d}}$. Here, $m = 6lbs$ is the mass of the hexapod, $g = 9.81m/s^2$ is the gravitational acceleration, $\rho = 1.225kg/m^2$ is the density of air, $A = 0.171m^2$ is the area of the base of the hexapod, and $C_d = 2.2916$ is the drag coefficient which results in a $V_t = 9.04m/s$. For this particular configuration, we find that the worst case speed of descent is in fact far less than the nominal speeds of autonomous vehicles on highways and freeways. In Figure 6, we plot the terminal velocity for objects of varying masses and areas, and observe that designs with lower mass and larger areas would incur fewer challenges in terms of motion blur and stability.

[1] https://www.bestprintingonline.com/resolution.htm

[2] https://www.amazon.com/Freenove-Big-Hexapod-Robot-Kit-Raspberry-Pi-Balancing-Recognition-Ultrasonic/dp/B08M5DXS2P

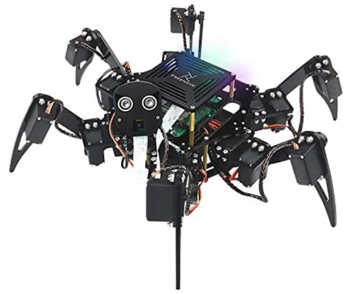

Figure 5: An example of a low-cost ground robot suitable for air-dropping. The robot weights approximately 6 Lbs., contains a Raspberry Pi 4 onboard for computation, and an operation time of roughly 40 minutes.

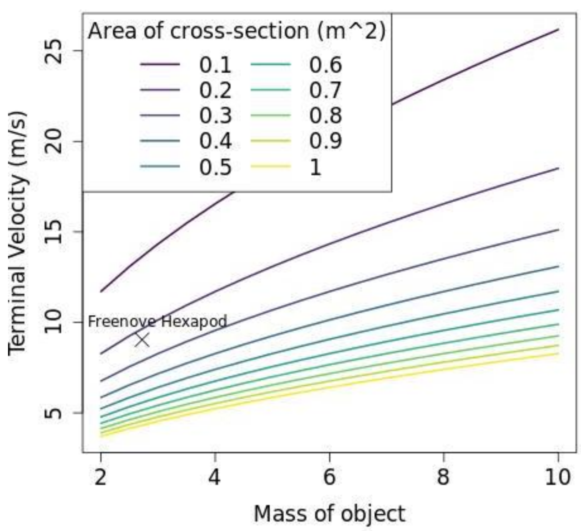

Figure 6: The impact of the mass (in kg) and cross-sectional area of the free falling object/robot on its terminal velocity.

5 RELATED WORK

Air–ground teaming: The use of unmanned autonomous systems (UAS) in performing aerial surveillance has been investigated for multitudes of scenarios (e.g., wildlife monitoring [3], illegal oil discharge [37], forest fires [31], etc.) owing to their capabilities in performing fast, broad-based monitoring. Recent works have also investigated the feasibility of cooperative networked aerial systems [8] addressing challenges related to decentralized control and planning. As Pippin et al. [33] describe, the heterogeneity in air and ground perspectives, i.e., the high coverage, albeit low resolution nature of aerial perspectives and the low coverage, high fidelity nature of ground-based sensing, is useful in several practical applications. Recently, Arora et al. [5] studied the problem of route planning for teams of unmanned aerial and ground vehicles considering the fuel constraints of UAVs, terrain constraints of UGVs and their speed differential, such that the UAV and UGV teammates rendezvouz for intermittent charging of the UAVs (using the equipment carried by the UGVs). Hood et al. [17] design a system where a UAV and UGV cooperate in exploring an unknown space (e.g., a damaged building) where the UAV guides the movement path of the UGV whereas the UGV maintains a fixed pose relative to the UAV by tracking and following a fiducial tag on the underside of the UAV. Deng at el. [11], recently, devised a sensor network for air–ground surveillance for improved target tracking. Garzon et al. [13] describe their early efforts in combining surveillance output from mini quadrotor and a ground robot to visualize obstacles and targets of interest. While these works describe efforts pertaining to various scenarios of how air–ground heterogeneous perspectives can be fused, they do not address terrain-aware navigation for the ground robots. Furthermore, these efforts typically involve (i) a single aerial system flying at a low altitude where issues of poor resolution are not dominant, (ii) and are designed for environments where active flight is a non-issue unlike our use cases where the regions are contested and minimizing the flight duration as well as the RF signatures that they produce are critical. The closest to our vision, Peterson et al. [32] describe their recent efforts in leveraging aerial perspectives for better terrain-awareness in ground robots. To account for low resolution inherent in images from higher altitudes, they propose having the UAS traverse set paths to cover larger ground at varying altitudes. Contrary to our proposed setting, this approach, however, would not be feasible in contested environments.

Light-weight Collaborative Perception at the Edge: Collaborative perception in static and mobile ground settngs, complementary to our vision, have been advocated recently [2, 18, 21, 35]. Qiu et al.[35] describe a scenario where cameras of differing capabilities co-exist in the same network: fixed surveillance cameras and resource-constrained mobile devices with cameras. The authors demonstrate that moving vehicles can be tracked seamlessly across this heterogeneous camera network through selective actuation of devices without overly draining the mobile devices. In essence, the resource-intensive video analytics pipeline is performed on the cloud and the mobile cameras are consulted intermittently, only to resolve ambiguities. Further, Lee et al. [21] demonstrate significant savings in bandwidth needs (of *dumb* cameras that offload raw footage to a central cloud) – they show that by establishing space-time relationships, apriori, between co-existing cameras, that they can be selectively turned ON (and OFF) leading to as much as 238 times savings in bandwidth at a miss rate of only 15% for a vehicle detection task. Similarly, Jain et al. [18] also show that significant correlations exist between co-located cameras, and discuss configurations of video analytics pipelines that can be triggered by peer cameras leading to both cost efficiency and superior inference accuracy. Unlike such past work, we focus explicitly on using collaboration to modify or abort the inferencing pipeline itself, instead of selectively activating nodes or performing fusion of the outputs from multiple nodes. Most recently, Yao et al. [2] describe the vision for providing machine intelligence as a service at the edge for resource-constrained devices. In addition to outlining core capabilities required for enabling such a service (e.g., scheduling, caching, resource profiling), they also describe opportunities for the convergence of the idea of collaboration between devices and deep intelligence as a service. Various recent works have described efforts in enabling lightweight collaboration ([1, 7, 19, 29] through different configurations (e.g., techniques that do not require re-training of the DNNs, and explore concepts such as region prioritization and partial execution of DNNs). The paradigm we propose in this paper will extend techniques from the ground domain to address aerial–to–ground challenges as well

6 CONCLUDING REMARKS

In this position paper, we propose a novel paradigm in air–ground teaming where spatially distributed, air-delivered robots collaborate in a resource-preserving manner to attain better reasoning over the terrain they will eventually navigate. Our preliminary evaluations show how lower resolution imaging (akin to those taken from very high altitudes, for e.g., by hovering drones or satellite imagery) can be detrimental to downstream applications such as object detection, and motivate the need for higher resolution, closer to ground aerial sensing. We discuss design goals and technical challenges ahead for building *AirDrop*.

ACKNOWLEDGEMENT

We acknowledge the support of the U.S. Army Grant No. W911NF21-20076.

REFERENCES

[1] T. Abdelzaher, Y. Hao, K. Jayarajah, A. Misra, P. Skarin, S. Yao, D. Weerakoon, and K.-E. Årzén. Five challenges in cloud-enabled intelligence and control. *ACM TOIT*, 20(1):1–19, 2020.

[2] T. Abdelzaher, S. Yao, and et al. Eugene: Towards deep intelligence as a service. In *Proc. of IEEE ICDCS'19*, 2019.

[3] K. Adams, A. Broad, D. Ruiz-García, and A. R. Davis. Continuous wildlife monitoring using blimps as an aerial platform: a case study observing marine megafauna. *Australian Zoologist*, 40(3):407–415, 2020.

[4] R. C. Arkin and R. R. Murphy. Autonomous navigation in a manufacturing environment. *IEEE Transactions on Robotics and Automation*, 6(4):445–454, 1990.

[5] D. Arora, P. Maini, P. Pinacho-Davidson, and C. Blum. Route planning for cooperative air-ground robots with fuel constraints: an approach based on cmsa. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 207–214, 2019.

[6] A. Bagula, L. Castelli, and M. Zennaro. On the design of smart parking networks in the smart cities: An optimal sensor placement model. *Sensors*, 15(7):15443–15467, 2015.

[7] H. Barahouei P and T. Nadeem. Collaborative intelligent cross-camera video analytics at edge: Opportunities and challenges. In *Proc. of AIChallengeIoT'19*, 2019.

[8] R. W. Beard, T. W. McLain, D. B. Nelson, D. Kingston, and D. Johanson. Decentralized cooperative aerial surveillance using fixed-wing miniature uavs. *Proceedings of the IEEE*, 94(7):1306–1324, 2006.

[9] F. Bonin-Font, A. Ortiz, and G. Oliver. Visual navigation for mobile robots: A survey. *Journal of intelligent and robotic systems*, 53(3):263–296, 2008.

[10] C. C. Castello, J. Fan, A. Davari, and R.-X. Chen. Optimal sensor placement strategy for environmental monitoring using wireless sensor networks. In *2010 42nd Southeastern Symposium on System Theory (SSST)*, pages 275–279. IEEE, 2010.

[11] X. Deng, Y. Liu, C. Zhu, and H. Zhang. Air–ground surveillance sensor network based on edge computing for target tracking. *Computer Communications*, 166:254–261, 2021.

[12] J. Ferryman and A. Shahrokni. Pets2009: Dataset and challenge. In *2009 Twelfth IEEE international workshop on performance evaluation of tracking and surveillance*, pages 1–6. IEEE, 2009.

[13] M. Garzón, J. Valente, D. Zapata, and A. Barrientos. An aerial–ground robotic system for navigation and obstacle mapping in large outdoor areas. *Sensors*, 13(1):1247–1267, 2013.

[14] P. Golej, L. Orlikova, J. Horak, P. Linhartova, and J. Struhar. Detection of people and vehicles using very high-resolution satellite images. *International Multidisciplinary Scientific GeoConference: SGEM*, 21(2.1):283–291, 2021.

[15] T. Guan, D. Kothandaraman, R. Chandra, A. J. Sathyamoorthy, K. Weerakoon, and D. Manocha. Ga-nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments. *IEEE Robotics and Automation Letters*, 7(3):8138–8145, 2022.

[16] Y. Guo, A. Song, J. Bao, T. Hongru, and J. Cui. A combination of terrain prediction and correction for search and rescue robot autonomous navigation. *International Journal of Advanced Robotic Systems*, 6(3):24, 2009.

[17] S. Hood, K. Benson, P. Hamod, D. Madison, J. M. O'Kane, and I. Rekleitis. Bird's eye view: Cooperative exploration by ugv and uav. In *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 247–255. IEEE, 2017.

[18] S. Jain, G. Ananthanarayanan, J. Jiang, Y. Shu, and J. Gonzalez. Scaling video analytics systems to large camera deployments. In *Proc. of ACM HotMobile*, 2019.

[19] K. Jayarajah, D. Wanniarachchige, T. Abdelzaher, and A. Misra. Comai: Enabling lightweight, collaborative intelligence by retrofitting vision dnns. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*, pages 41–50. IEEE, 2022.

[20] V. Khryashchev and R. Larionov. Wildfire segmentation on satellite images using deep learning. In *2020 Moscow Workshop on Electronic and Networking Technologies (MWENT)*, pages 1–5. IEEE, 2020.

[21] J. Lee and a. et. On tracking realistic targets in a megacity with contested air and spectrum access. *MILCOM'18*.

[22] S. Liu, X. Fu, M. Wigness, P. David, S. Yao, L. Sha, and T. Abdelzaher. Self-cueing real-time attention scheduling in criticality-aware visual machine perception. In *2022 IEEE 28th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 173–186. IEEE, 2022.

[23] S. Liu, S. Yao, X. Fu, R. Tabish, S. Yu, A. Bansal, H. Yun, L. Sha, and T. Abdelzaher. On removing algorithmic priority inversion from mission-critical machine inference pipelines. In *2020 IEEE Real-Time Systems Symposium (RTSS)*, pages 319–332. IEEE, 2020.

[24] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.

[25] Z. Liu, Z. Li, B. Liu, X. Fu, I. Raptis, and K. Ren. Rise of mini-drones: Applications and issues. In *Proceedings of the 2015 workshop on privacy-aware mobile computing*, pages 7–12, 2015.

[26] V. Lukoševičius, R. Keršys, A. Keršys, R. Makaras, and J. Jablonskytė. Three-and four-mass models for vehicle front crumple zone. *Transport Problems*, 15, 2020.

[27] J. G. Masek. Stability of boreal forest stands during recent climate change: evidence from landsat satellite imagery. *Journal of biogeography*, 28(8):967–976, 2001.

[28] G. Mateo-Garcia, J. Veitch-Michaelis, L. Smith, S. V. Oprea, G. Schumann, Y. Gal, A. G. Baydin, and D. Backes. Towards global flood mapping onboard low cost satellites with machine learning. *Scientific reports*, 11(1):1–12, 2021.

[29] A. Misra, K. Jayarajah, D. Weerakoon, R. Tandriansyah, S. Yao, and T. Abdelzaher. Dependable machine intelligence at the tactical edge. In *SPIE'19*.

[30] P. Nguyen, H. Truong, M. Ravindranathan, A. Nguyen, R. Han, and T. Vu. Matthan: Drone presence detection by identifying physical signatures in the drone's rf communication. In *Proceedings of the 15th annual international conference on mobile systems, applications, and services*, pages 211–224, 2017.

[31] A. Ollero and L. Merino. Unmanned aerial vehicles as tools for forest-fire fighting. *Forest Ecology and Management*, 234(1):S263, 2006.

[32] J. Peterson, H. Chaudhry, K. Abdelatty, J. Bird, and K. Kochersberger. Online aerial terrain mapping for ground robot navigation. *Sensors*, 18(2):630, 2018.

[33] C. Pippin, G. Gray, M. Matthews, D. Price, A.-P. Hu, W. Lee, M. Novitzky, and P. Varnell. The design of an air-ground research platform for cooperative surveillance. Technical report, Technical Report. Tech. Rep. 112010, Georgia Tech Research Institute, 2010.

[34] H. Qiu, P. Huang, N. Asavisanu, X. Liu, K. Psounis, and R. Govindan. Autocast: Scalable infrastructure-less cooperative perception for distributed collaborative driving. *arXiv preprint arXiv:2112.14947*, 2021.

[35] H. Qiu, X. Liu, S. Rallapalli, A. J. Bency, K. Chan, R. Urgaonkar, B. Manjunath, and R. Govindan. Kestrel: Video analytics for augmented multi-camera vehicle tracking. In *Proc. of IEEE/ACM IoTDI'18*. IEEE, 2018.

[36] A. J. Sathyamoorthy, K. Weerakoon, T. Guan, J. Liang, and D. Manocha. Terrapn: Unstructured terrain navigation through online self-supervised learning. *arXiv preprint arXiv:2202.12873*, 2022.

[37] N. Serra-Sogas, P. D. O'Hara, R. Canessa, P. Keller, and R. Pelot. Visualization of spatial patterns and temporal trends for aerial surveillance of illegal oil discharges in western canadian marine waters. *Marine pollution bulletin*, 56(5):825–833, 2008.

[38] S. Skakun, N. Kussul, A. Shelestov, and O. Kussul. Flood hazard and flood risk assessment using a time series of satellite images: A case study in namibia. *Risk Analysis*, 34(8):1521–1537, 2014.

[39] K. Weerakoon, A. J. Sathyamoorthy, U. Patel, and D. Manocha. Terp: Reliable planning in uneven outdoor environments using deep reinforcement learning. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 9447–9453. IEEE, 2022.

[40] S. Weidling, N. Todtenberg, T. Basmer, M. Schoelzel, M. Maaser, and M. Taubert. Outdoor range measurements in sub-ghz license-free radio bands under realistic conditions. In *Proceedings of the 15th ACM International Symposium on Mobility Management and Wireless Access*, pages 71–74, 2017.