From Passive Observation to Active Multi-Agent Sensing in Earth Observation

Xiaojian Wang* Zhengxiong Li*
*University of Colorado Denver
{xiaojian.wang, zhengxiong.li}@ucdenver.edu

Abstract

Multimodal Earth observation has framed the "imperfect lens" of cloud cover, resolution shortfall, asynchronous radar/optical acquisitions, and irregular revisit as a modeling *problem, addressed with fusion, imputation, cross-modal translation, and foundation-model pretraining. We argue these imperfections are also a* sampling *problem. Implicit in this framing is an assumption that observation is* taken as given*: revisit schedules are set by orbit, pre/post-event pairs are whatever was collected, and ground truth is annotated long after the fact. What is missing is a sensing layer that is now widely deployed in the world: autonomous robotic platforms such as aerial robots, ground robots, and autonomous vessels, which can be* tasked *on the basis of what satellites cannot see. We position robots as active sampling agents in a multimodal Earth observation loop: they consume satellite-derived uncertainty maps as a tasking signal, contribute high-resolution observations as a modality alongside satellite data, and provide* in-situ *ground truth that closes the loop on online model adaptation. We diagnose four mismatches between current passive-observation practice and what active multi-agent observation requires, and trace three concrete design moves toward an Earth observation stack in which satellites and robots act as co-observers rather than separate data pipelines.*

1. Introduction

Multimodal Earth observation (EO) has matured into a sophisticated machinery for extracting Earth-process information from heterogeneous, asynchronous, often-incomplete satellite data. Foundation models pretrained on billions of pixels can fine-tune to flood mapping, burn-scar segmentation, building damage assessment, and crop monitoring [4, 9, 22]. Methods that fuse synthetic-aperture radar (SAR) and optical imagery reconstruct optical-like representations through clouds [13]. Multimodal benchmarks now bracket disaster events with pre-event optical and post-event SAR imagery, reflecting realistic operational constraints [2, 6].

Underneath this progress sits an assumption inherited from the satellite-only era: that observation is taken as given. Whatever the satellite's revisit schedule and cloud mask deliver becomes the input; the modeling task is to extract as much signal as possible from that fixed input. This framing is reasonable for the satellite layer alone. It has also shaped the way much of the community thinks about "imperfect" observation: imperfection tends to be cast as a *modeling* problem, and the methods that emerge (e.g., temporal imputation, cross-modal translation, cloud removal, super-resolution) are largely responses to the question *what do we do with the data we happened to receive?*

The satellite layer gives EO a broad view of the planet, but it also fixes the terms of observation: the orbit, the revisit time, the cloud mask, and the sensor payload. Robotic platforms complement this view by making part of the observing process taskable. They do not observe everywhere, but they can be directed toward the places where the broad view becomes uncertain, incomplete, or operationally urgent. In this role, robots are not just additional sensors; they are *controllable* sensors. Autonomous robotic platforms now operate at scale in exactly the environments EO most cares about: search-and-rescue ground robots in earthquake zones [16, 18], fixed-wing unmanned aerial vehicle (UAV) teams over wildfire perimeters [10], agricultural sensing platforms across cellular dead zones [20], autonomous surface vessels in coastal waters. Where they go, when they observe, what fidelity they collect, and which targets they prioritize are decision variables. Robotics has spent decades developing the corresponding theory of active perception, informative path planning, and active simultaneous localization and mapping (SLAM) [1, 7, 14]. Yet multimodal EO rarely integrates robotic platforms as a co-observation modality within the same inference loop as satellite data.

We argue that this separation is more a matter of convention than necessity: as robotic platforms become widely deployed, the opportunity to integrate them into the satellite observation loop deserves closer examination. Beyond serving as another modality, robotic platforms can act as an

This work is partially supported by US NSF Awards #2426470 and #2436428.

Table 1. Four mismatches between current multimodal EO practice and the potential of active multi-agent observation.

Current passive-observation practice	What robotic platforms make possible	What an active multi-agent loop enables
Satellite revisit time is fixed by orbit; pre/post-event pairs are whatever was collected [2, 6].	Robotic platforms (aerial, ground, surface, and underwater) can be tasked at minute-to-hour revisit times in the same scene [10, 18].	Trigger an observation during the critical window, instead of waiting for the next satellite pass.
Cloud cover and resolution shortfalls are treated as a modeling problem, addressed with imputation and cross-modal fusion [13].	A UAV can fly below the cloud, a ground robot can approach the relevant building, and an underwater vehicle can observe below the surface.	Send a robot to physically reach what the satellite cannot resolve, instead of modeling around the gap.
Foundation models [4, 9, 22] are pre-trained offline; downstream fine-tuning relies on annotated benchmarks [2, 6].	Robots in the field generate *in-situ* observations of damage, vegetation, and infrastructure during the actual event.	Update the satellite-side model online with in-situ robot observations, instead of waiting for offline annotation.
Existing multimodal EO benchmarks [2, 6] are built from satellite imagery and manual annotations; robotic observations remain rare as a primary input modality.	Robotic observations are themselves a sensing modality with distinctive characteristics such as high resolution, sparse coverage, and irregular timing [3, 5, 12].	Integrate robot observations alongside satellite data as a full modality in the EO pipeline.

active sampling layer whose tasking can be conditioned on what the satellite layer cannot resolve. Conversely, satellite-derived uncertainty maps can serve as a tasking signal that current robotic mission planners do not consume.

Position. We propose designing Earth observation as an *active multi-agent observation system*. Satellites provide broad but temporally and physically constrained background coverage, while robots provide targeted, on-demand observations in the gaps; together, they coordinate by sharing what each side needs and what each side can supply.

Contributions. We (1) diagnose four mismatches between current passive-observation practice and what active multi-agent observation requires (§2); (2) state the position and identify the three properties that distinguish it (§3); (3) trace three concrete design moves: uncertainty-driven robot tasking, robot observations as a fusion modality, and closed-loop *in-situ* model adaptation (§4). We close with discussion (§5), related work (§6), and conclusion (§7).

2. The Passive-Observation Bias

To motivate the position, we first examine where multimodal EO practice implicitly assumes that observation is fixed. Table 1 summarizes four such mismatches between current practice and the alternatives that an active sampling layer would unlock.

We highlight three observations from the table. First, each row is a place where multimodal EO has invested heavily on the modeling side (imputation, fusion, domain adaptation, foundation-model pretraining), with much less attention paid to the sampling side. These investments are valuable, but they share a common starting point: *they take the data as given*. Second, each alternative builds on robotic platforms already operating in the field (Sec. 1); among the open challenges is designing the interface that connects them to the satellite-side EO model. Third, the mismatches reinforce each other. A model that does not know it could request a robotic observation cannot identify the locations where one would help; a robot mission planner that does not consume satellite uncertainty cannot prioritize them. Closing any one mismatch in isolation tends to capture only part of the available value.

3. Position: Active Multi-Agent Observation

The mismatches in Table 1 share a common cause: much of multimodal EO has inherited the satellite-only abstraction, in which the observer is fixed and the world is sampled by orbit. That abstraction becomes increasingly limiting as autonomous robotic platforms are deployed alongside satellites. Robotic platforms have become widely available, mission-driven, and controllable, and a mature body of work in active perception, informative path planning, and active SLAM [1, 7, 14] provides the algorithmic foundation for treating sampling itself as a decision variable. This layer has yet to be systematically integrated into multimodal EO.

Our position:

> *Earth observation can be designed as an active multi-agent system: satellites provide broad coverage constrained by physical and operational factors; robotic platforms supply targeted sampling where satellite observation is insufficient; and a coordination layer connects the two.*

Fig. 1 sketches the loop this position implies. Three properties distinguish it. First, **sampling is a planning variable**. Current multimodal EO places its decision-making downstream of observation: given the data, process it. The active position places decision-making at both

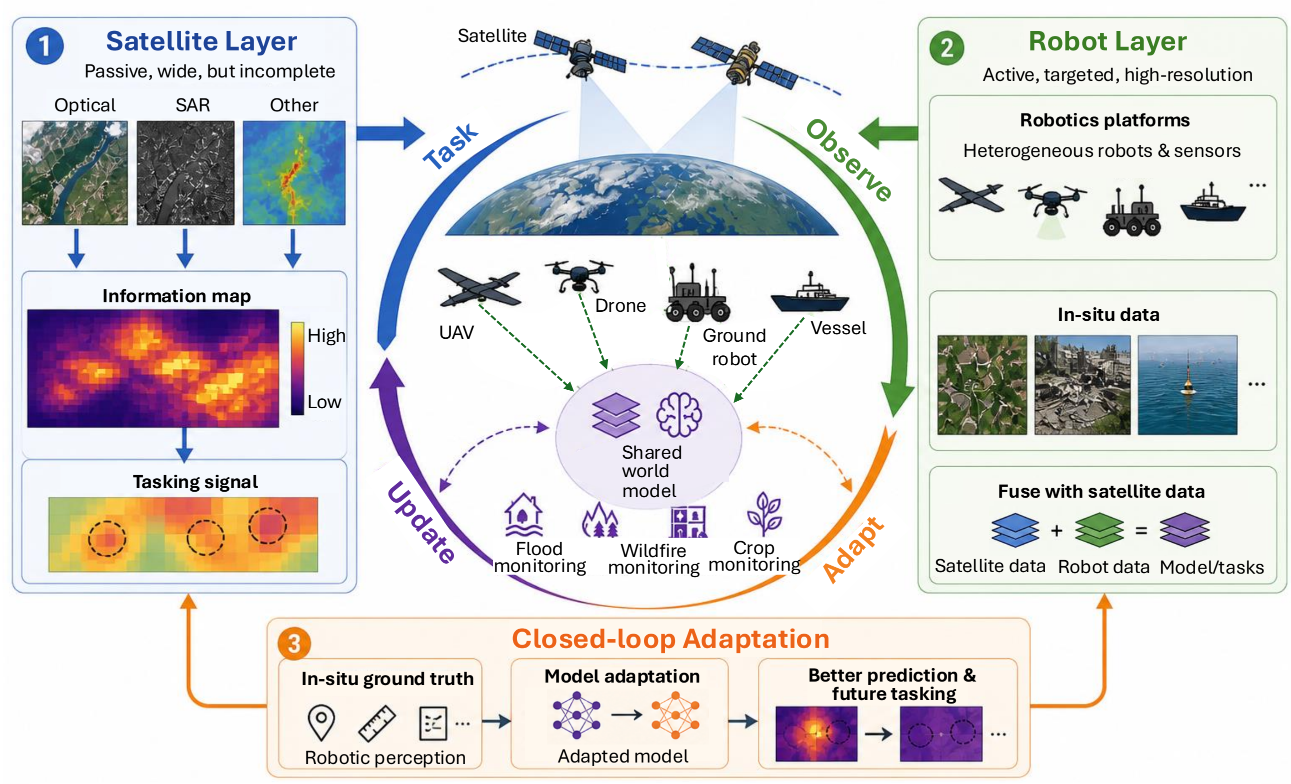

Figure 1. Active multi-agent Earth observation loop. The satellite layer (①) provides broad but imperfect coverage and emits an uncertainty/information map that serves as a tasking signal. Robotic platforms (②), such as UAVs, ground robots, and autonomous vessels, consume this signal, collect targeted high-resolution and *in-situ* observations, and return them as a modality fused with satellite data. A shared world model and monitoring objectives at the center coordinate both layers. Closed-loop adaptation (③) uses *in-situ* ground truth to adapt the EO model, yielding better predictions and refined future tasking.

ends: given a forecast of what we need to know, also decide what to observe next. A natural source for that forecast is the satellite-side EO model's predictive uncertainty. Second, **uncertainty is the coupling**. Satellite-side uncertainty (cloud cover, resolution shortfall, model disagreement, last-good-observation age) tells the robot where to look. Robot-side constraints (coverage, energy, mission) tell the satellite what the robot can actually deliver. The two layers can communicate primarily through uncertainty rather than through raw pixels. Third, **closing the loop is increasingly within reach**. Robotic platforms equipped with modern middleware [11] and cloud-and-fog robotics stacks [8] provide the building blocks for online interaction with EO models. Low Earth Orbit (LEO) infrastructure underpins both ends of the loop: imaging constellations supply the short revisit times needed for fresh satellite-side observation, and communication constellations supply the low-latency links needed to reach the field. The two have complementary strengths: satellites provide broad spatial coverage while robots provide targeted local detail. The position concerns *how they should be designed to interact*.

4. Three Design Moves

We trace three concrete design moves the position implies, organized by where the active multi-agent loop is closed: at the robot's planner (§4.1), at the multimodal model's fusion layer (§4.2), and at its online learning loop (§4.3). The three reinforce each other. Without uncertainty-driven tasking, the robot's data may not target the model's actual gaps. Without modality-level fusion, robotic observations tend to remain outside the EO pipeline, serving mainly as a source of labels or validation. Without online adaptation, the model tends to remain at its pretrained state, regardless of what the robot collects.

4.1. Robot Tasking Conditioned on EO Uncertainty

Today's robotic mission planners typically consume environment models, sensor coverage maps, and mission objectives, and plan trajectories largely against those. These inputs typically do not include the satellite-side EO model's predictive uncertainty: per-pixel cloud probability, model-confidence maps, time-since-last-good-observation, ensemble disagreement, damage-class entropy. The first design move is to expose these as a tasking signal that the robot's planner consumes alongside its mission cost.

A wildfire-monitoring UAV [10] should bias its trajectory toward regions where the satellite-derived burn perimeter is most uncertain, not just toward the nominal fire front. A damage-assessment UAV deployed after a hurricane [2] should prioritize neighborhoods where post-event SAR-only coverage produces low-confidence building-level classification, because optical ground truth at those locations is what the model lacks. Informative path planning provides the algorithmic substrate for trading off mission utility against information gain [14]; what is missing is an interface in which the EO model's uncertainty is the information signal. The move requires a compact representation of EO uncertainty that can be transmitted to fielded platforms, a tasking objective that balances mission utility with uncertainty reduction, and a cooperation pattern that allows EO model providers to share uncertainty estimates without exposing model internals.

4.2. Robot Observations as a Modality

Multimodal fusion in EO has matured along axes that fit regular-revisit satellite sensors: optical $\times$ SAR [2, 13], multispectral $\times$ multispectral, optical $\times$ light detection and ranging (LiDAR). Robot-collected observations do not always satisfy the implicit regularities those architectures assume. A UAV-flown sub-cloud optical patch may cover only a small footprint at much higher resolution than the satellite, last only briefly during the critical window, and align with no satellite acquisition. A ground-robot photograph may be taken from below the rooftop the satellite sees from above. These observations constitute a sensing modality with distinctive characteristics that can be integrated into the EO pipeline alongside satellite data.

The second design move is to extend multimodal EO architectures to admit robot-collected observations as a full modality integrated into the EO pipeline. Foundation models for EO [4, 9] are typically built on assumptions of fixed channel sets and regular spatial layouts; admitting robot observations requires architectures that handle sparse, irregular, variable-fidelity inputs at inference. Concretely, the move requires representations that ingest observations defined by their pose and time rather than a fixed grid, fusion mechanisms that gracefully degrade when a robot modality is absent, and benchmarks [2] extended with robotic observations such as UAV imagery or ground-level photographs alongside existing satellite data.

4.3. Closed-Loop *In-Situ* Adaptation

Foundation models for EO are pretrained offline, often on data drawn from disaster events distinct from the one currently in progress. Distribution shift can be substantial: new construction styles, novel disaster types, and unobserved regions tend to degrade off-the-shelf models [2]. A common mitigation is post-hoc fine-tuning once labels are eventually annotated. The active position turns this around: robots in the field generate *in-situ* observations of the actual event, in real time, providing ground-level evidence that a satellite-only model otherwise lacks.

The third design move is to close the model-adaptation loop with the active sampling layer. As robots collect observations guided by the model's uncertainty (§4.1), those observations can in turn be used to adapt the model, creating a feedback loop between sensing and learning. The move requires online adaptation methods stable under robot-prioritized streams that are not independent and identically distributed (non-IID), protocols for handling active-learning selection bias, and safeguards against adversarial or sensor-failure inputs.

5. Discussion

Connectivity and deployability. The active loop assumes that robots and EO model providers can exchange uncertainty maps, observations, and model updates over the duration of a mission. In many disaster-affected dead zones, this is also where terrestrial connectivity tends to be missing. Emerging LEO satellite networks have the potential to close this gap, but the bandwidth, latency, and operator-side cooperation required for closed-loop adaptation remain open. In the meantime, partial implementations (e.g., offline batching of robot observations, periodic re-tasking) deliver some of the value without continuous connectivity.

Heterogeneity of robotic platforms. Robotic platforms span aerial, ground, surface, and underwater systems, and vary widely in coverage, sensing, and energy budget. A UAV may survey a large area quickly but carry limited payload, while a ground robot captures fine-grained detail but moves slowly over a small area. Concrete instantiations will need a tasking framework that abstracts over diverse platform capabilities and allocates tasks according to each platform's strengths and constraints.

Evaluation. Existing multimodal EO benchmarks [2, 6] are built from pre-collected satellite observations with post-hoc annotations. Evaluating the position would require benchmarks that capture closed-loop behavior, where each observation influences subsequent model updates and tasking decisions. Designing such benchmarks (e.g., through replay, simulation, or pre-defined sampling policies) is an important open challenge for the community.

Failure modes and graceful degradation. The active loop has multiple potential points of failure: a robot sensor can return bad data, communication links can drop, the EO model can drift, and corrupted observations can destabilize online adaptation. A practical implementation needs to degrade gracefully when components fail, falling back on whatever signals remain available. Standards for partial-loop operation, drift monitoring, and rollback remain open.

Cost and triggering economics. Robot operations carry real cost in energy, time, airspace, and operator attention. An active loop that triggered a robot every time the model was uncertain would be hard to scale in practice. Tasking objectives that balance mission utility, expected information gain, and operational cost will become increasingly important as the system scales.

Privacy. Active sampling raises privacy concerns at multiple levels. At the data level, robotic platforms operate in closer proximity to people and property than satellites, potentially capturing personally identifiable information such as faces or license plates. Data collection, retention, and sharing policies must account for this. At the model level, shared uncertainty maps or model updates exchanged between satellites and robots may inadvertently reveal sensitive information about monitored regions or populations. Addressing these concerns will require privacy-aware protocols across the entire active observation loop.

6. Related Work

Multimodal EO and foundation models. A growing literature builds models that fuse heterogeneous satellite modalities and pretrains on global, multi-temporal data collections [4, 9, 13, 22]. Disaster-response benchmarks pair pre-event and post-event imagery, often crossing modalities to handle cloud cover and acquisition constraints [2, 6]. A complementary position treats satellite data as a distinct modality [15]. Our work extends this view by adding an *active* sensing layer: robots tasked on the basis of what satellites cannot resolve.

Active perception, informative path planning, and active SLAM. The robotics community has decades of work on *why*, *when*, and *where* an embodied agent should observe [1], with mature algorithmic foundations for trading off mission utility against information gain [7] and for exploration that reduces map uncertainty [14]. This work takes the agent's own internal uncertainty as the information signal; we propose extending it to *external* uncertainty supplied by an EO model.

Active learning and adaptive sampling in EO. Active learning frames *what to label next* as a function of model state, and the remote sensing community has long applied it to satellite image classification [19]. Our position takes a step further upstream: rather than choosing which existing pixel to annotate, the active multi-agent loop chooses which new pixel to physically acquire. The information-theoretic objectives developed for label selection transfer, but the cost model changes from annotation effort to robot mission cost.

Test-time and online adaptation. A growing body of work updates a model on the inference distribution it actually encounters, either through self-supervised auxiliary losses at test time [17] or by minimizing prediction entropy on a target stream [21]. These methods typically assume the test stream arrives passively. Our third design move (§4.3) couples test-time adaptation to an active sampler: the stream is shaped by what the robot is tasked to observe, which is in turn shaped by the model's current state.

UAV and robotic remote sensing. A substantial remote sensing literature documents unmanned aerial systems for fine-grained environmental monitoring across photogrammetry, agriculture, and ecology [3, 12], and the broader robotics community has surveyed mobile platforms for in-situ environmental sensing [5]. This line of work has largely focused on the platforms and the data they produce; integrating those outputs into the multimodal EO modeling stack alongside satellite data remains relatively unexplored.

Disaster-response robotics and connectivity. Field-deployed robotic systems (ground, aerial, and surface) demonstrate operation at scale in disaster environments [10, 16, 18]. Cloud-and-fog robotics platforms [8] on top of Robot Operating System 2 [11] bring cloud-scale models within reach of fielded robots, and rural Internet of Things (IoT) platforms [20] demonstrate connectivity solutions for remote sensing in resource-limited environments. These efforts focus on coordination and connectivity among robots; the EO model typically remains a downstream consumer rather than an active participant in the loop. We propose integrating the EO model as an active component in the loop, both consuming and shaping robotic observations.

7. Conclusion

Multimodal Earth observation has matured into a sophisticated discipline for extracting signal from imperfect satellite data, with cloud cover, resolution shortfall, and asynchronous acquisition all addressed through ever-better models. We have argued that the data imperfections commonly addressed through modeling admit a complementary, sampling-side response: rather than only modeling around the gaps, the field can also choose what gets observed in the first place. A second sensing layer of autonomous robotic platforms is widely deployed today and can be tasked on the basis of what the satellite layer cannot resolve. Three design moves follow: tasking robots from satellite-side uncertainty, fusing their observations as a modality at inference, and closing the model-adaptation loop with the in-situ ground truth they collect. Both layers have matured largely on their own; exploring how to bring them into a single loop is an open and promising direction. This agenda spans benchmarks, coordination protocols between satellite and robotic systems, robust online adaptation, and privacy-aware frameworks. The growing maturity of foundation-model EO, robotic autonomy, and connectivity infrastructure suggests that progress on these fronts is increasingly feasible. A multimodal EO that observes the world actively, rather than waiting for it, is a goal worth pursuing.

References

[1] Ruzena Bajcsy, Yiannis Aloimonos, and John K Tsotsos. Revisiting active perception. *Autonomous Robots*, 42(2):177–196, 2018. 1, 2, 5

[2] Hongruixuan Chen, Jian Song, Olivier Dietrich, Clifford Broni-Bediako, Weihao Xuan, Junjue Wang, Xinlei Shao, Yimin Wei, Junshi Xia, Cuiling Lan, et al. Bright: A globally distributed multimodal building damage assessment dataset with very-high-resolution for all-weather disaster response. *Earth System Science Data*, 17(11):6217–6253, 2025. 1, 2, 4, 5

[3] Ismael Colomina and Pere Molina. Unmanned aerial systems for photogrammetry and remote sensing: A review. *ISPRS Journal of photogrammetry and remote sensing*, 92:79–97, 2014. 2, 5

[4] Yezhen Cong, Samar Khanna, Chenlin Meng, Patrick Liu, Erik Rozi, Yutong He, Marshall Burke, David Lobell, and Stefano Ermon. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. *Advances in Neural Information Processing Systems*, 35:197–211, 2022. 1, 2, 4, 5

[5] Matthew Dunbabin and Lino Marques. Robots for environmental monitoring: Significant advancements and applications. *IEEE Robotics & Automation Magazine*, 19(1):24–39, 2012. 2, 5

[6] Ritwik Gupta, Bryce Goodman, Nirav Patel, Ricky Hosfelt, Sandra Sajeev, Eric Heim, Jigar Doshi, Keane Lucas, Howie Choset, and Matthew Gaston. Creating xbd: A dataset for assessing building damage from satellite imagery. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 10–17, 2019. 1, 2, 4, 5

[7] Geoffrey A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014. 1, 2, 5

[8] Jeffrey Ichnowski, Kaiyuan Chen, Karthik Dharmarajan, Simeon Adebola, Michael Danielczuk, Víctor Mayoral-Vilches, Nikhil Jha, Hugo Zhan, Edith LLontop, Derek Xu, et al. Fogros2: An adaptive platform for cloud and fog robotics using ros 2. In *2023 IEEE international conference on robotics and automation (ICRA)*, pages 5493–5500. IEEE, 2023. 3, 5

[9] Johannes Jakubik, Sujit Roy, CE Phillips, Paolo Fraccaro, Denys Godwin, Bianca Zadrozny, Daniela Szwarcman, Carlos Gomes, Gabby Nyirjesy, Blair Edwards, et al. Foundation models for generalist geospatial artificial intelligence. *arXiv preprint arXiv:2310.18660*, 2023. 1, 2, 4, 5

[10] Kyle D Julian and Mykel J Kochenderfer. Distributed wildfire surveillance with autonomous aircraft using deep reinforcement learning. *Journal of Guidance, Control, and Dynamics*, 42(8):1768–1778, 2019. 1, 2, 4, 5

[11] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. Robot operating system 2: Design, architecture, and uses in the wild. *Science robotics*, 7(66):eabm6074, 2022. 3, 5

[12] Salvatore Manfreda, Matthew F McCabe, Pauline E Miller, Richard Lucas, Victor Pajuelo Madrigal, Giorgos Mallinis, Eyal Ben Dor, David Helman, Lyndon Estes, Giuseppe Ciraolo, et al. On the use of unmanned aerial systems for environmental monitoring. *Remote sensing*, 10(4):641, 2018. 2, 5

[13] Andrea Meraner, Patrick Ebel, Xiao Xiang Zhu, and Michael Schmitt. Cloud removal in sentinel-2 imagery using a deep residual neural network and sar-optical data fusion. *ISPRS Journal of Photogrammetry and Remote Sensing*, 166:333–346, 2020. 1, 2, 4, 5

[14] Julio A Placed, Jared Strader, Henry Carrillo, Nikolay Atanasov, Vadim Indelman, Luca Carlone, and José A Castellanos. A survey on active simultaneous localization and mapping: State of the art and new frontiers. *IEEE Transactions on Robotics*, 39(3):1686–1705, 2023. 1, 2, 4, 5

[15] Esther Rolf, Konstantin Klemmer, Caleb Robinson, and Hannah Kerner. Position: Mission critical–satellite data is a distinct modality in machine learning. In *Forty-first International Conference on Machine Learning*, 2024. 5

[16] Maira Saboia, Lillian Clark, Vivek Thangavelu, Jeffrey A Edlund, Kyohei Otsu, Gustavo J Correa, Vivek Shankar Varadharajan, Angel Santamaria-Navarro, Thomas Touma, Amanda Bouman, et al. Achord: Communication-aware multi-robot coordination with intermittent connectivity. *IEEE Robotics and Automation Letters*, 7(4):10184–10191, 2022. 1, 5

[17] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pages 9229–9248. PMLR, 2020. 5

[18] Marco Tranzatto, Takahiro Miki, Mihir Dharmadhikari, Lukas Bernreiter, Mihir Kulkarni, Frank Mascarich, Olov Andersson, Shehryar Khattak, Marco Hutter, Roland Siegwart, et al. Cerberus in the darpa subterranean challenge. *Science Robotics*, 7(66):eabp9742, 2022. 1, 2, 5

[19] Devis Tuia, Michele Volpi, Loris Copa, Mikhail Kanevski, and Jordi Munoz-Mari. A survey of active learning algorithms for supervised remote sensing image classification. *IEEE Journal of Selected Topics in Signal Processing*, 5(3):606–617, 2011. 5

[20] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta Sinha, Ashish Kapoor, Madhusudhan Sudarshan, and Sean Stratman. {FarmBeats}: an {IoT} platform for {Data-Driven} agriculture. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 515–529, 2017. 1, 5

[21] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020. 5

[22] Zhitong Xiong, Yi Wang, Fahong Zhang, Adam J Stewart, Joëlle Hanna, Damian Borth, Ioannis Papoutsis, Bertrand Le Saux, Gustau Camps-Valls, and Xiao Xiang Zhu. Neural plasticity-inspired multimodal foundation model for earth observation. *arXiv preprint arXiv:2403.15356*, 2024. 1, 2, 5