Appears in Proceedings of the 14th International Workshop on Planning and Scheduling for Space, Toulouse, France, April 2025.

# Leveraging Realtime Meteorological Data for Dynamic Tasking of Agile Earth-Observing Satellites

**Shreeyam Kacker [1], Steve Chien [2], Kiruthika Devaraj [3], Işıl Demir [3], Kerri Cahoy [1]**
[1]Massachusetts Institute of Technology
[2]Jet Propulsion Laboratory, California Institute of Technology
[3]Planet Labs
shreeyam@mit.edu

## Abstract

Observing unpredictable phenomena is challenging for Earth observation missions due to long lead times from scheduling, uplinking, and executing the image capture onboard the spacecraft. This delay means that conditions can change in between planning and execution. Dynamic tasking is a mission concept that aims to mitigate this unpredictability by moving autonomy onboard the spacecraft and quickly reacting to conditions as observed, with multiple potential perception sources that can be used to inform decision making for observation. Dynamic tasking has been proposed using a variety of information sources: lookahead sensors, other satellites, ground sensors, and airborne and marine sources.

Modern meteorological satellites in Geostationary Earth Orbit (GEO) are designed to provide fast refresh rates for imagery for weather monitoring and prediction, with up to 10 minute full-disk image cadence at 20 minutes latency after L1 processing. Imagery from meteorological satellites could potentially be used as virtual lookahead sensors to update target utility for Earth-observing satellites, from which the imaging schedule can be updated either onboard or via a ground station.

In this work, we explore the usage of global meteorological data from GOES West, GOES East, Meteosat Zero Degree Service (ZDS), Meteosat Indian Ocean Data Coverage (IODC) and Himawari for the application of dynamic tasking for cloud avoidance for agile Earth-observing satellites. As there are many methods to uplink information to a spacecraft, such as through ground stations, data relay satellite networks, and direct reception of meteorological data, assessing overall latency is not possible without a specific architecture. We instead parameterize latency between 0 to 2 hours of uplink delay, both through a real-time link and through intermittent contact through ground stations. We then compare dynamically tasked scheduling performance against both conventional (non-dynamically tasked) methods and a theoretical omniscient scheduler to establish bounds for the results using a dataset of world cities as a proxy for global satellite imagery demand.

## Introduction

Earth-observing satellites are typically split into "monitoring" and "tasked" missions. Monitoring missions aim to capture as much imagery as possible without having specific targets, giving rapid response and broad area coverage, typically at lower spatial resolutions due to constraints on aperture size, sensor pitch, and data budgets (King et al. 2013; NASA 2024). In contrast, tasked missions follow predetermined schedules to capture specific targets at higher spatial resolutions, but require external input to direct the tasking (Planet Labs PBC 2017; Maxar Technologies 2025). Both monitoring and tasked missions struggle to capture unpredictable phenomena: monitoring missions may require large amounts of imagery in order to capture a single event, while a tasked missions may make perform suboptimal observations if the event is no longer present when imagery is captured.

Unpredictable phenomena can either serve as attractors of tasks (high-priority events to be imaged) or detractors (undesirable conditions affecting imaging tasks). The most common detractor in Earth observation, aside from in meteorological studies, is cloud cover, which at any given time obscures approximately two-thirds of Earth's surface (King et al. 2013). Task attractors include applications such as real-time wildfire tracking (Nolde et al. 2021; Lenzen et al. 2014), deep convective ice storm studies (Swope et al. 2024), and planetary boundary layer research (Candela et al. 2024).

To address these challenges, dynamic targeting has been proposed as a mission concept from NASA (Chien and Troesch 2015; Candela, Swope, and Chien 2023; Swope et al. 2024), CNES (Damiani, Verfaillie, and Charmeau 2005; Beaumet, Verfaillie, and Charmeau 2011), and is in operation by JAXA on GOSAT-2 (Suto et al. 2021; Imasu et al. 2023). While the specific implementation of dynamic tasking can vary, the overall goal is to move more autonomy onboard the spacecraft to re-assess conditions dynamically, and re-optimize imaging activities in accordance with updated information in a single overflight of a target. Recent advancements in onboard computing have enabled the feasibility of real-time decision-making using modern vision systems and onboard scheduling algorithms (Planet Labs PBC 2024; Kacker et al. 2022a,b; Kacker and Cahoy 2024). Sensor webs have also been previously used for rapid targeting of phenomena such as volcano eruptions (Chien et al. 2020), floods (Chien et al. 2019) and other dynamic phenomena. In this work, we use dynamic tasking to refer to dynamic targeting, but with an initial schedule solution.

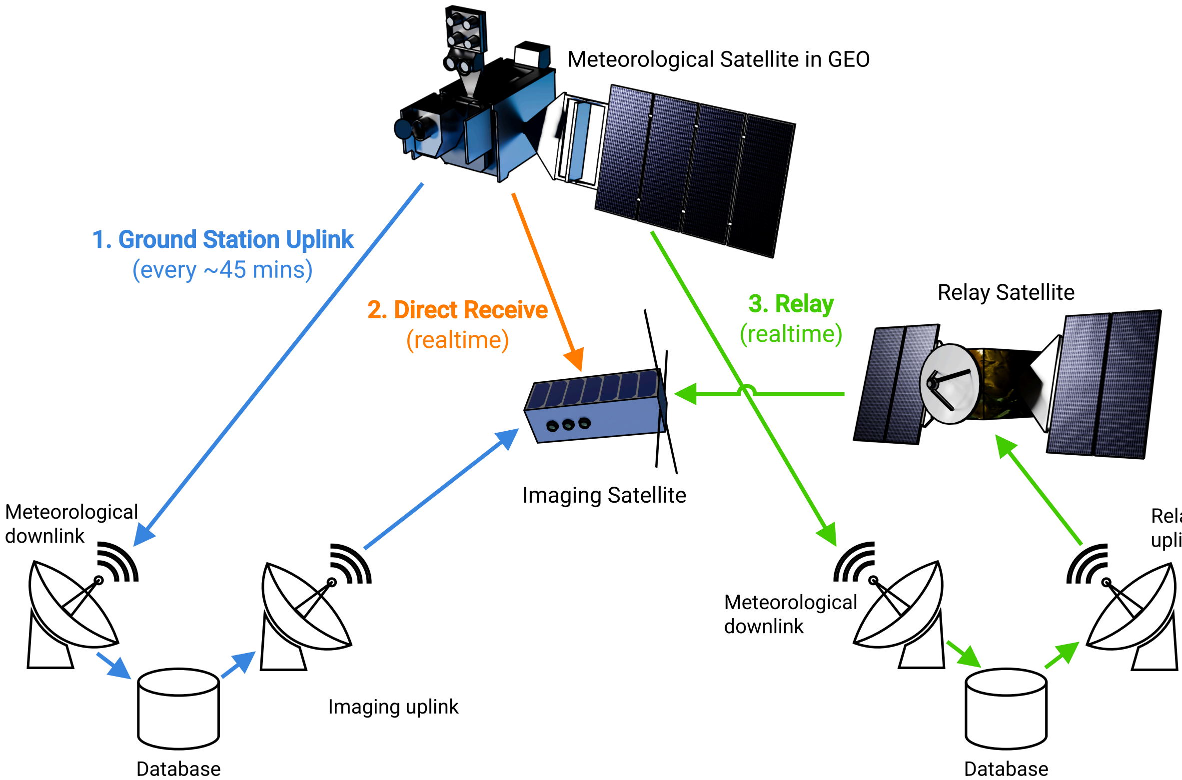


Figure 1: Diagram showing methods by which data from meteorological satellites can be transferred to imaging satellites. (1) Data can be relayed through downlinking, extracting from a database, and uplinking to the satellite, resulting in communication gaps. (2) Data can also be obtained in near realtime through direct reception by the imaging satellite through an antenna, or (3) relayed through a network such as Iridium, GlobalStar, or the NASA Tracking and Data Relay Satellite (TDRS) system.

## Background

### Geostationary Meteorological Satellites

Meteorological satellites in GEO provide crucial observations, especially over uninhabited areas and over oceans to inform weather monitoring. Spacecraft from many different nations cover different regions of the Earth. These spacecraft are typically equipped with high resolution, high refresh rate scanning instruments with many different bands. The scanning instruments on these spacecraft can deliver full-disk images up to a rate of every 10 minutes, with specific regions of Earth being scanned even faster than that to allow for tracking of fast-moving weather events (Mahonchak 2019). The specific service regions, satellites, instruments, center longitudes, and full-disk image cadences used in this work are provided in Table 1.

There are three main ways that data can be processed and transferred from a meteorological satellite to an imaging satellite. Images from the meteorological satellites are downlinked through a high-rate radio, that allows for L1 imagery to be available through storage systems such as through Amazon S3 20 minutes after capture (Mahonchak 2019), with processed L2 imagery being processed and delivered later. Since the data is publicly available, it can be retrieved by the imaging satellite's ground station, and then either (1) uplinked directly, or (3) relayed by another satellite through a system like Iridium, GlobalStar, or the NASA Tracking and Data Relay Satellite (TDRS) system. Additionally, through low-rate channels such as through GOES' High Rate Information Transmission / Emergency Managers Weather Information Network (HRIT/EMWIN) low-rate broadcast system, masks can also be (2) directly downlinked to the imaging satellite, although this requires additional licensing and hardware (D'Anzeo et al. 2025). Using either the relay system or the direct downlink system allows for real-time decision making, whereas communication gaps from a ground station means updates can only be transferred approximately every 45 minutes for an imaging satellite in LEO. An overview of all the methods and their cadence is shown in Fig. 1.

In this work, we use the L2 binary cloud mask product to evaluate whether imaging activites are cloudy or not, although recent advancements in cloud masking algorithms (Thompson et al. 2014; Wagstaff et al. 2018) allow for creating cloud masks directly from L1 or L0 data.

We consider data from the Geostationary Operational Environmental Satellites (GOES) program, the Meteosat Second Generation (MSG) program, and the Himawari program. Fig. 2 shows the total theoretical coverage of these spacecraft along with the largest 10,000 world cities used

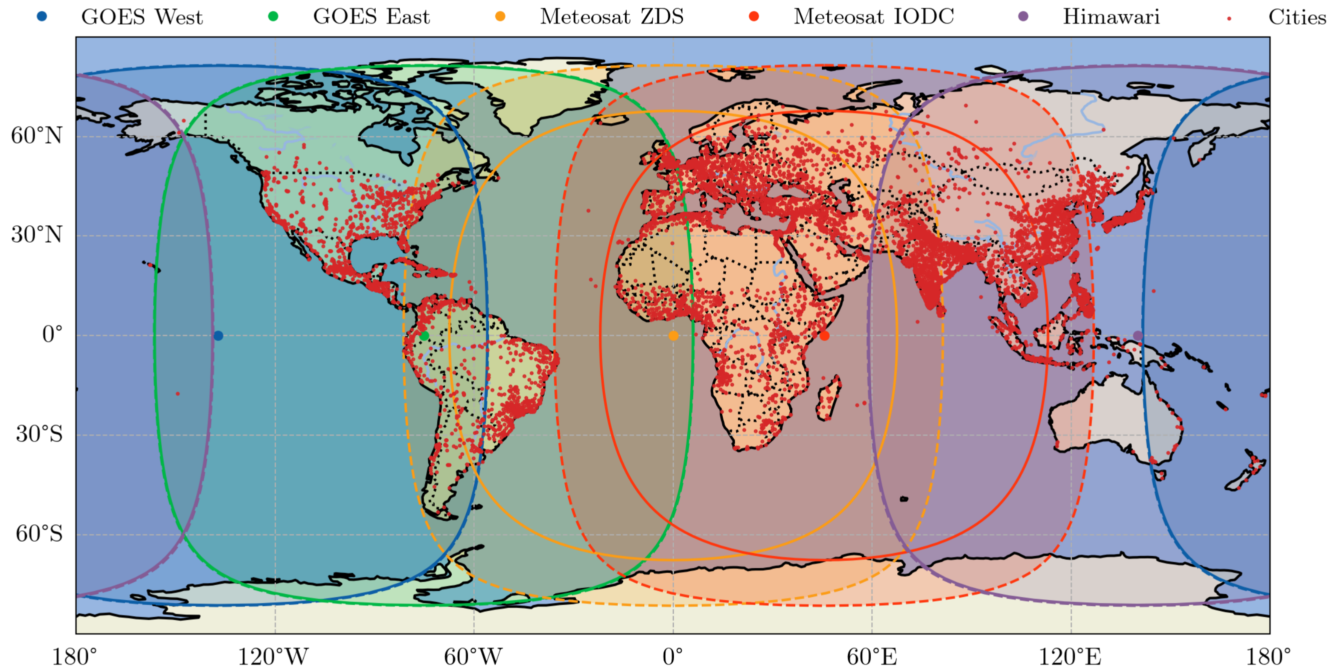


Figure 2: Theoretical (dotted) and actual (solid) coverage of selected meteorological satellites. Actual data coverage of products such as cloud masks is typically narrower, due to complexities of deriving reflectance data at grazing incidences. Largest 10,000 world cities used as an imaging dataset are also overlaid (Simplemaps 2010).

in this work as a qualitiative proxy for global satellite imagery demand (Simplemaps 2010; Planet Labs PBC 2017). While the 10,000 world cities dataset has intrinsic bias, it is qualitatively a good fit for the commercial applications targeted in this work, and the analysis can easily be extended to consider at other factors. Fig. 2 shows that the theoretical coverage of these spacecraft spans the globe, up to approximately 80 degrees latitude, although with ground sample distance (GSD) trailing dramatically at grazing incidences, starting at 1 km for most bands at nadir and 10+ km at imaging limits. L2 data such as binarized cloud masks however are typically more limited in coverage as compared to L1 data, due to the complexity of obtaining reflectance values for grazing incidences. Cloud masking algorithms from these spacecraft can rely on emissive bands as well as reflective bands, hence cloud masks can be obtained even at night, allowing for global coverage, although with artifacts along the day/night terminator line due to rapidly changing irradiance values from the sun.

### Dynamic Tasking

The Autonomous Sciencecraft Experiment on Earth Observing One (Chien et al. 2005) pioneered onboard re-planning and machine learned classifiers (Castano et al. 2006) but lacked computational power for single-overflight planning, focusing instead on follow-up observations. More modern approaches have been able to leverage massively increased compute power onboard spacecraft (Rijlaarsdam et al. 2024; Kacker et al. 2022a) and extend to constellations (Gorr et al. 2025).

Currently, the most advanced dynamic tasking system can be found on JAXA's GOSAT-2 satellite, which utilizes a re-orientable fold mirror and onboard compute system to screen for clouds and point to the least cloudy area, which results in a 150-250% increase in usable imagery (Nassar et al. 2023; Imasu et al. 2023; Suto et al. 2021; Oishi et al. 2017).

## Approach

Our approach is split up into two parts: First, we evaluate the predictive ability of global cloud masks obtained from satellites for acquiring data at a later time, and then we incorporate data from those cloud masks into a dynamic tasking simulation, re-optimizing the schedule based on cloud cover information obtained from the cloud masks.

### Cloud Masks as Priors

To optimize planning horizons and quantify timescales over which imagery from meteorological satellites is useful, we evaluate the predictive power of cloud masks by using them as a prediction. Given data $D$ and a prior cloud mask prior from an observation from an earlier time $M_{\text{prior}}$, we can compare this to the naive flat prior, assuming all tasks are cloud free. Given that clouds cover approximately 66% of the Earth's surface at any one time (King et al. 2013), the accuracy of this naive prediction $M_{\text{naive}}$ will be fixed around 34%.

From these two models of the data, we can then calculate

Table 1: Selected meteorological satellites in GEO providing full-disk imagery, with associated instrument, center longitude, and full-disk image cadence, current as of date of publication.

| **Service** | **Satellite** | **Instrument** | **Center Longitude** | **Full-Disk Image Cadence** |
|---|---|---|---|---|
| GOES East | GOES-16 | ABI | 75.2° W | 10 minutes |
| GOES West | GOES-18 | ABI | 137.2° W | 10 minutes |
| Meteosat ZDS | Meteosat-10 | SEVIRI | 0° E | 15 minutes |
| MeteoSat IODC | Meteosat-9 | SEVIRI | 45.5° E | 15 minutes |
| Himawari | Himawari-9 | AHI | 140.7° E | 10 minutes |

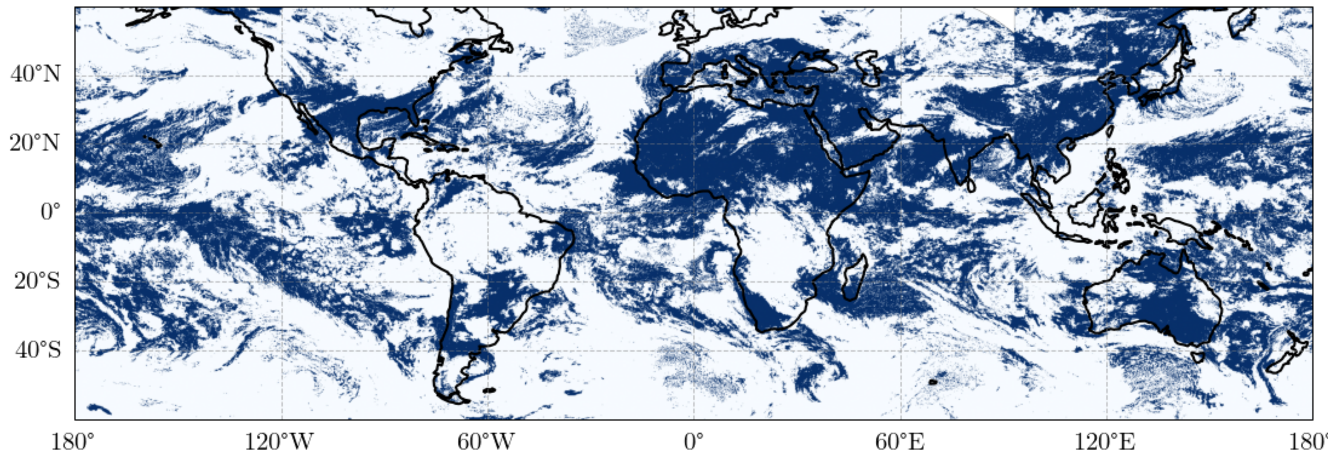


Figure 3: Example binary cloud mask from stitching observations from selected meteorological spacecraft at time 2025-01-01T00:00:00+00. White areas denote cloudy regions, blue areas denote non-cloudy regions, and coastlines are highlighted in black.

the Bayes factor in order to evaluate at a high level the predictive ability of using a previous cloud mask as a prior for later masks, over time, as compared to the naive model. The Bayes factor is given by

$$\mathrm{BF} = \frac{P(D \mid M_{\text{prior}})}{P(D \mid M_{\text{naive}})}. \tag{1}$$

## Scheduling System

We then utilize prior cloud masks as data to inform our dynamic tasking algorithm. We simulate a scheduling system with world cities as targets, agility constraints on spacecraft imaging, and a utility model where imaging requests that are occluded by clouds have zero utility, unity otherwise. Formally, this can be described as a constraint satisfaction problem, and as a further subset, can be represented as a mixed-integer linear program (MILP). We use a MILP formulation for the scheduling problem, same as in (Augenstein et al. 2016; Nag et al. 2019; Nag, Li, and Merrick 2018; Eddy 2021). The MILP can be defined as:

$$\begin{aligned} \text{maximize} \quad & \sum_{x_i \in X} f(x_i)\, x_i \\ \text{subject to} \quad & \forall i,\ x_i \in \{0,1\} \\ & x_i + x_j \leq 1\ \forall k(x_i, x_j) = 0, \end{aligned}$$

where $X$ is the set of imaging accesses from a set of imaging requests $R$, where $R$ in our case is the dataset of world cities. Since a particular imaging request may have multiple points in the scheduling horizon where it can be captured, the mapping from requests to imaging accesses is non-unique. $f$ is a function evaluating the utility of an access (equal to unity in this work), and $k$ represents a set of constraints. In this work we primarily assume **agility** and **repetition** constraints, which are conventionally the tightest constraints for spacecraft scheduling (Eddy and Kochenderfer 2020; Augenstein 2014).

Agility constraints consist of the requirement to slew from one task to another along the cross-track of the spacecraft. We use a model for agility that assumes constant acceleration and constant deceleration for slews without any roll rate limits, along with a fixed duration settling time. Hence,

$$t_{\text{slew}} = t_s + \alpha\sqrt{|\Delta\theta|}, \tag{2}$$

Table 2: Summary of all parameters used in dynamic tasking simulation.

| Parameter | Symbol | Value |
|---|---|---|
| Schedule Block Start | | 2025-01-01T00:06:00+00 |
| Schedule Block End | | 2025-01-01T00:18:00+00 |
| Schedule Horizon | | 12 hours |
| Orbit Inclination | $i$ | 51.6° |
| Orbit Altitude | $h$ | 400 km |
| Orbit Period | $T$ | 5553.5 s |
| Field of Regard | | 30° |
| Settling Time | $t_s$ | 10 s |
| Slewing Acceleration | $\alpha$ | 4.0 s/$\sqrt{\text{deg}}$ |

where $t_s$ is the spacecraft settling time, and $\Delta\theta$ is the angle difference between two tasks, with $\alpha$ a parameter representing slewing acceleration of the spacecraft. The constraints themselves can be given by:

$$k_{\text{agility}}\ (x_i, x_j) = \begin{cases} 1 & \text{if } t_{\text{slew}}(\theta_j - \theta_i) \leq t_j - t_i \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

requiring that all slews must take less time than the maneuver and settling time of the spacecraft. Finally, the final constraint is repetition,

$$k_{\text{repetition}}\ (x_i, x_j) = \begin{cases} 1 & \text{if } r(x_i) = r(x_j) \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

where $r(x)$ returns the request $r \in R$ corresponding to an imaging access.

Additionally, we consider a single spacecraft in ISS orbit and only take tasking opportunities during non-eclipse phases of orbit. A full table of parameters used for the scheduler simulation is given in Table 2.

## Results

Fig. 4 shows the probability of a later binarized cloud mask being equal to a prior. The probability drops off rapidly, before asymptoting the expected values of approximately 66% cloud cover. Fig. 5 shows the Bayes factor of using a previous cloud mask as a prior up to 35.5 hours, compared to the naive prior used in conventional satellite scheduling where all tasks are assumed cloud-free. Initially the mask-based prior has significantly more predictive power than the naive, before dropping down and asymptoting at around twice the power. This behavior is likely due to individual clouds being hard to predict, but their formation conditions e.g. near mountains being relatively stagnant, resulting in aggregate higher predictive power.

Fig. 6 shows the same data, but now as a locus on a precision-recall plot. With the assumption that baseline predictive power is achieved after 35.5 hours reaching, precision and recall appear to decay linearly to an ultimate value of approximately (0.73, 0.77), showing no strong class bias. Predictive power also reaches half its baseline value after two hours, emphasizing the importance of timeliness in using prior masks.

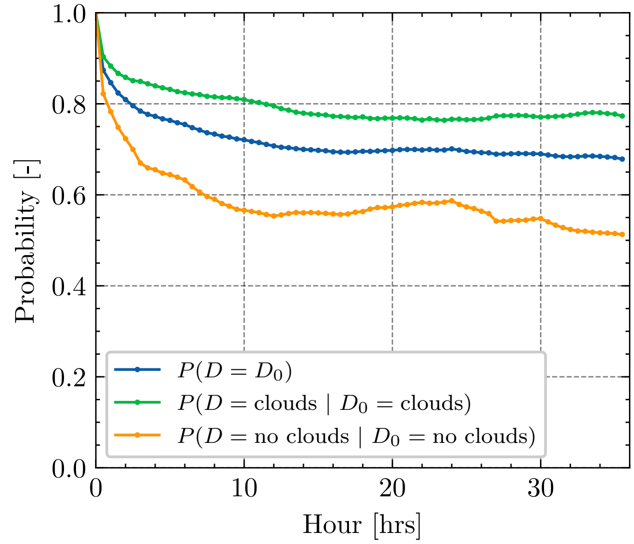


Figure 4: Probability of cloud status being predicted by a previous cloud mask from the past, split by total, and initial classification. Data is taken from a global cloud mask at 2025-01-01T00:00:00+00:00 looking up to 35 hours in the future.

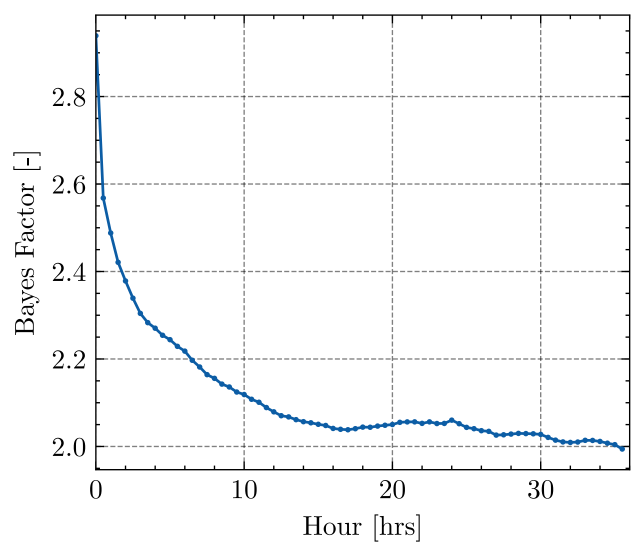


Figure 5: Bayes factor of comparing a prior equal to an earlier observation with a naive prior that all tasks are always cloud-free, showing significantly higher predictive ability of the earlier observation compared to the naive model. The Bayes factor initially drops off rapidly before asymptoting around a value of 2.0, showing that latency is a big factor in predictive ability. Data is taken from a global cloud mask at 2025-01-01T00:00:00+00:00 looking up to 35.5 hours in the future.

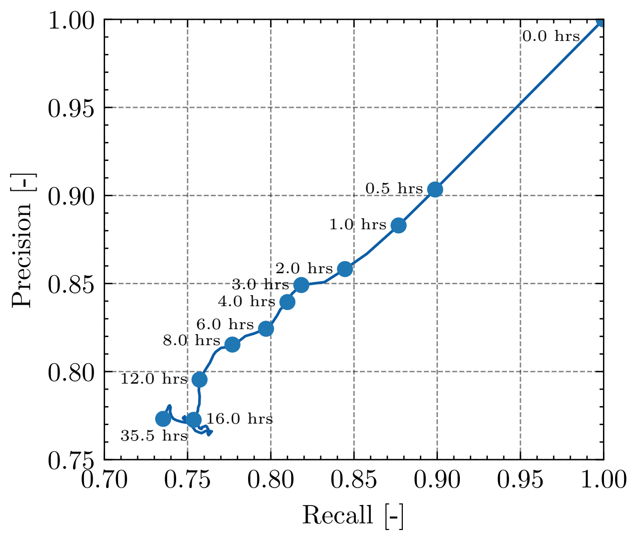


Figure 6: Precision-recall curve for cloud prediction from prior masks with a time delay, showing diminishing predictive ability of cloud mask priors with time. Data is taken from a global cloud mask at 2025-01-01T00:00:00+00:00 looking up to 35.5 hours in the future. Within two hours, most of the predictive power compared to the baseline is lost.

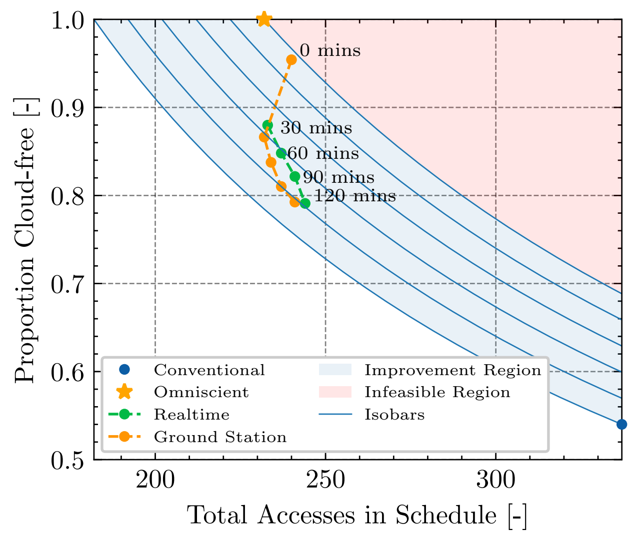


Figure 7: Conventional, omniscient, and delayed schedules compared in terms of total length of schedule, and proportion of captures cloud-free, showing that larger information delay leads to overall worse schedule performance, transiting perpendicular to the total cloud free imagery isobars. All delayed (realtime vs ground station access) schedules are approximately the same number of imaging accesses of the omniscient schedule, with decreasing proportion cloud-free as delay length increases. Schedules transmitted through ground stations perform about 2 percentage points worse than those with real-time communications, but a theoretical schedule with instant information transmitted over ground stations is within 5% of the omniscient solution in total number of cloud-free images, as shown by isobars representing total number of cloud-free images.

Fig. 7 then compares using different latencies for scheduling, with ground station based communication gaps and a real-time link, bounding between to a conventional and omniscient scheduler and highlighting between these with isobars representing total number of cloud-free captures. The conventional scheduler is essentially sampling from the stationary distribution of clouds, resulting in approximately half of its captures being cloudy as expected, due to favorable weather conditions along the schedule. The omniscient scheduler is a theoretical best-case bound, using a scheduler that has futuresight, and subtends an infeasibility region, as the MILP solver guarantees that the omniscient schedule is optimal in terms of total number of images captured for the agility and repetition constraints imposed.

With a 30 minute information latency, a schedule consisting of 89% of captures being cloud-free can be obtained, showing the power of using cloud masks as priors for scheduling and prediction. Ground station communication gaps result in approximately 2 percentage points performance penalty in cloud-free imagery, although roughly being on par with the real-time link after two hours. This effect is likely due to the average additional delay due to communication being approximately 22.5 minutes, which becomes more insignificant as overall data delay increases. A theoretical instant link with zero delay between capture, processing, and downlink to a ground station is also shown, which can capture over 95% of the improvement in total number of cloud-free imagery, given by its position within the isobars.

## Conclusions and Future Work

In conclusion, we show that virtual lookahead using meteorological data from satellites in GEO can be advantageous for Earth observation in cases where cloud cover is undesirable, but is sensitive to latency requirements and will require streamlined infrastructure to extract the most benefit. Real-time always-connected communication systems improve efficiency of cloud-free collects the most, by 63%, but are very close to the total number of useful collects transmitting information only through ground stations every half orbit at 60.4%, both with a 30 minute delay. As the data delay reduces, the gap between these two data points will likely increase, as the delay intrinsic in using a ground station for communications becomes more significant. Further investigation will be conducted on using the maximum imaging frequency of the meteorological satellite platforms at 10 and 15 minutes, but will require additional handling of data connectivity due to the mixed rates involved, with Meteosat operating at 15 minute cadence and all the other platforms considered operating at 10 minute cadence (Mahonchak 2019; EUMETSAT 2020).

Based on these results, short time horizon weather prediction from providers such as Tomorrow.io may also be beneficial for dynamic spacecraft operations, relaxing the requirements for streamlined, low-latency infrastructure or real-time links (Tomorrow.io 2025).

This work looks at aggregate task performance in a schedule and shows that even a two-hour delay can be beneficial for improving useful image throughput, but it is possible that using data from cloud mask priors results in poor predictive performance in certain geographic regions, which will require further investigation. Areas like the Amazon rainforest in particular may have poor predictive performance due to the dynamic nature of weather systems in the region. Along these lines, certain orbits and certain configurations of ground stations may also be better than others, which will require additional investigation.

The results in this work are taken looking forward from a single point in time. Sampling across different seasons would make these conclusions significantly more robust, as seasonality likely has an effect on predictive power.

## Acknowledgments

This work was generously funded by Planet Labs. We are grateful to Planet Labs for their collaboration.

Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). Government Sponsorship acknowledged.

## References

Augenstein, S. 2014. Optimal Scheduling of Earth-Imaging Satellites with Human Collaboration via Directed Acyclic Graphs. *AAAI Spring Symposium Series*.

Augenstein, S.; Estanislao, A.; Guere, E.; and Blaes, S. 2016. Optimal scheduling of a constellation of earth-imaging satellites, for maximal data throughput and efficient human management. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 26, 345–352.

Beaumet, G.; Verfaillie, G.; and Charmeau, M.-C. 2011. Feasibility of Autonomous Decision Making on Board an Agile Earth-Observing Salite. *Computational Intelligence*, 27(1): 123–139. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8640.2010.00375.x.

Candela, A.; Swope, J.; and Chien, S. A. 2023. Dynamic Targeting to Improve Earth Science Missions. *Journal of Aerospace Information Systems*, 20(11): 679–689.

Candela, A.; Victoria, J. D.; Zilberstein, I.; Kurowski, M.; Yue, Q.; and Chien, S. 2024. Dynamic Targeting Scenario to Study the Planetary Boundary Layer. *International Geoscience and Remote Sensing Symposium*.

Castano, R.; Mazzoni, D.; Tang, N.; Greeley, R.; Doggett, T.; Cichy, B.; Chien, S.; and Davies, A. 2006. Onboard classifiers for science event detection on a remote sensing spacecraft. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 845–851.

Chien, S.; Mclaren, D.; Doubleday, J.; Tran, D.; Tanpipat, V.; and Chitradon, R. 2019. Using Taskable Remote Sensing in a Sensor Web for Thailand Flood Monitoring. *Journal of Aerospace Information Systems (JAIS)*, 16(3): 107–119.

Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davis, A.; Mandl, D.; Frye, S.; Trout, B.; Shulman, S.; and Boyer, D. 2005. Using Autonomy Flight Software to Improve Science Return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication*, 2(4): 196–216.

Chien, S.; and Troesch, M. 2015. Heuristic Onboard Pointing Re-scheduling for an Earth Observing Spacecraft.

Chien, S. A.; Davies, A. G.; Doubleday, J.; Tran, D. Q.; Mclaren, D.; Chi, W.; and Maillard, A. 2020. Automated Volcano Monitoring Using Multiple Space and Ground Sensors. *Journal of Aerospace Information Systems (JAIS)*, 17:4: 214–228.

Damiani, S.; Verfaillie, G.; and Charmeau, M.-C. 2005. A Continuous Anytime Planning Module for an Autonomous Earth Watching Satellite. In *Proceedings of the 25th International Conference on Automated Planning and Scheduling*.

D'Anzeo, F.; Dell, T.; Rodegari, M.; Chien, S.; Candela, A.; Yue, Q.; Kurowski, M.; and Zilberstein, I. 2025. Passive interception, decoding and processing of public broadcast L-Band EO data in LEO. In *Proceedings of the 5th ESA Workshop on Advanced Flexible Telecom Payloads*. ESA/ECSAT Harwell Campus, Didcot, United Kingdom.

Eddy, D. 2021. *Task Planning for Earth Observing Satellite Systems*. Ph.D. thesis.

Eddy, D.; and Kochenderfer, M. 2020. Markov Decision Processes For Multi-Objective Satellite Task Planning. In *2020 IEEE Aerospace Conference*, 1–12. Big Sky, MT, USA: IEEE. ISBN 978-1-7281-2734-7.

EUMETSAT. 2020. Meteosat series. https://www.eumetsat.int/our-satellites/meteosat-series.

Gorr, B.; Jaramillo, A. A.; Gao, H.; Selva, D.; Mehta, A.; Sun, Y.; Ravindra, V.; David, C. H.; and Allen, G. H. 2025. Decentralized Satellite Constellation Replanning for Event Observation. *Journal of Spacecraft and Rockets*, 0(0): 1–19. Publisher: American Institute of Aeronautics and Astronautics _eprint: https://arc.aiaa.org/doi/pdf/10.2514/1.A36143.

Imasu, R.; Matsunaga, T.; Nakajima, M.; Yoshida, Y.; Shiomi, K.; Morino, I.; Saitoh, N.; Niwa, Y.; Someya, Y.; Oishi, Y.; Hashimoto, M.; Noda, H.; Hikosaka, K.; Uchino, O.; Maksyutov, S.; Takagi, H.; Ishida, H.; Nakajima, T. Y.; Nakajima, T.; and Shi, C. 2023. Greenhouse gases Observing SATellite 2 (GOSAT-2): mission overview. *Progress in Earth and Planetary Science*, 10(1): 33.

Kacker, S.; and Cahoy, K. 2024. Reinforcement-Learned Lookahead Heuristics for Earth-Observing Satellites. *Small Satellite Conference*.

Kacker, S.; Meredith, A.; Cahoy, K.; and Labrèche, G. 2022a. Machine Learning Image Processing Algorithms Onboard OPS-SAT. *Small Satellite Conference*.

Kacker, S.; Meredith, A.; Kusters, J.; Tomio, H.; Felt, V.; and Cahoy, K. 2022b. On-orbit rule-based and deep learning image segmentation strategies. *AIAA SCITECH 2022 Forum*. MAG ID: 4205910275 S2ID: 6647e3fa789a96d068eb556b2d346155a0b0eb52.

King, M. D.; Platnick, S.; Menzel, W. P.; Ackerman, S. A.; and Hubanks, P. A. 2013. Spatial and Temporal Distribution of Clouds Observed by MODIS Onboard the Terra and Aqua Satellites. *IEEE Transactions on Geoscience and Remote Sensing*, 51(7): 3826–3852. Conference Name: IEEE Transactions on Geoscience and Remote Sensing.

Lenzen, C.; Woerle, M. T.; Göttfert, T.; Mrowka, F.; and Wickler, M. 2014. Onboard Planning and Scheduling Autonomy within the Scope of the FireBird Mission. In *SpaceOps 2014 Conference*. Pasadena, CA: American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-221-9.

Mahonchak, G. D. 2019. GOES-R Series Data Book.

Maxar Technologies. 2025. WorldView Legion | Maxar.

Nag, S.; Li, A.; Alan S. Li; Alan S. Li; Vinay Ravindra; Ravindra, V.; Vinay Ravindra; Marc Sanchez Net; Net, M. S.; Cheung, K.-M.; Lammers, R.; and Bledsoe, B. P. 2019. Autonomous Scheduling of Agile Spacecraft Constellations with Delay Tolerant Networking for Reactive Imaging. MAG ID: 2998771969.

Nag, S.; Li, A. S.; and Merrick, J. H. 2018. Scheduling algorithms for rapid imaging using agile Cubesat constellations. *Advances in Space Research*, 61(3): 891–913.

NASA. 2024. VIIRS | NASA Earthdata. Publisher: Earth Science Data Systems, NASA.

Nassar, R.; MacDonald, C. G.; Kuwahara, B.; Fogal, A.; Issa, J.; Girmenia, A.; Khan, S.; and Sioris, C. E. 2023. Intelligent pointing increases the fraction of cloud-free CO2 and CH4 observations from space. *Frontiers in Remote Sensing*, 4: 1233803.

Nolde, M.; Plank, S.; Richter, R.; Klein, D.; and Riedlinger, T. 2021. The DLR FireBIRD Small Satellite Mission: Evaluation of Infrared Data for Wildfire Assessment. *Remote Sensing*, 13(8): 1459. Number: 8 Publisher: Multidisciplinary Digital Publishing Institute.

Oishi, Y.; Ishida, H.; Nakajima, T.; Nakamura, R.; and Matsunaga, T. 2017. The Impact of Different Support Vectors on GOSAT-2 CAI-2 L2 Cloud Discrimination. *Remote Sensing*, 9(12): 1236.

Planet Labs PBC. 2017. Hi-Res SkySat Imagery Now Available.

Planet Labs PBC. 2024. Planet Labs PBC Announces Real-Time Insights Technology Using NVIDIA Jetson Platform.

Rijlaarsdam, D.; Hendrix, T.; González, P. T. T.; Velasco-Mata, A.; Buckley, L.; Miquel, J. P.; Casaled, O. A.; and Dunne, A. 2024. The Next Era for Earth Observation Spacecraft: An Overview of CogniSAT-6.

Simplemaps. 2010. World Cities Database | Simplemaps.com.

Suto, H.; Kataoka, F.; Kikuchi, N.; Knuteson, R. O.; Butz, A.; Haun, M.; Buijs, H.; Shiomi, K.; Imai, H.; and Kuze, A. 2021. Thermal and near-infrared sensor for carbon observation Fourier transform spectrometer-2 (TANSO-FTS-2) on the Greenhouse gases Observing SATellite-2 (GOSAT-2) during its first year in orbit. *Atmospheric Measurement Techniques*, 14(3): 2013–2039. Publisher: Copernicus GmbH.

Swope, J.; Chien, S.; Bosch-Lluis, X.; Dunkel, E.; Yue, Q.; Ogut, M.; Ramos, I.; Kangaslahti, P.; Deal, W.; and Cooke, C. 2024. Storm Classification and Dynamic Targeting for a Smart Ice Cloud Sensing Satellite. *Journal of Aerospace Information Systems*, 21(12): 958–971.

Thompson, D. R.; Green, R. O.; Keymeulen, D.; Lundeen, S. K.; Mouradi, Y.; Nunes, D. C.; Castaño, R.; and Chien, S. A. 2014. Rapid Spectral Cloud Screening Onboard Aircraft and Spacecraft. *IEEE Transactions on Geoscience and Remote Sensing*, 52(11): 6779–6792. Conference Name: IEEE Transactions on Geoscience and Remote Sensing.

Tomorrow.io. 2025. Tomorrow.io.

Wagstaff, K.; Chien, S.; Altinok, A.; Rebbapragada, U.; Thompson, D.; Schaffer, S.; and Tran, D. 2018. Cloud Filtering and Novelty Detection using Onboard Machine Learning for the EO-1 Spacecraft. In *International Symposium on Artificial Intelligence, Robotics, and Automation for Space (ISAIRAS 2018)*. Madrid, Spain. Also appears at AI in the Oceans and Space Workshop, International Joint Conference on Artificial Intelligence (IJCAI 2017).