EXPANDING AND MATURING DYNAMIC TARGETING

Steve Chien[1], **Alberto Candela**[1], **Juan Delfa**[1], **Akseli Kangaslahti**[1,2], **and Abigail Breitfeld**[1,3]

[1]*Jet Propulsion Laboratory,[*] California Institute of Technology, USA*
[2]*University of Michigan, USA*
[3]*Carnegie Mellon University, USA*

ABSTRACT

Dynamic Targeting (DT) is a mission concept in which data from a lookhead sensor is continuously analyzed to determine how to target a primary sensor. We review the base concept and discuss several efforts to mature the concept. First we discuss efforts to extend the DT concept to more complex slewing and observation utility models as well as discuss realistic execution timing. Second, we describe investigations to learn DT targeting policies. Third, we discuss flight of DT on the ESA OPS SAT testbed.

Key words: Artificial Intelligence, Space Applications, flight software, space mission operations.

1. INTRODUCTION

Dynamic Targeting (DT) is a mission concept in which instrument observations are used to rapidly reconfigure and repoint an instrument to optimize science observations.

Studies have already shown DT to be a promising technology for cloud avoidance and hunting deep convective ice storms. In the cloud avoidance use case, a lookahead instrument is used to identify clouds allowing an agile imager to avoid identified clouds to acquire cloud free targets. In the storm hunting case a lookahead sensor is used to identify plausible storm centers allowing an electronically steered radar to target to increase the yield of deep convective storms.

We also describe challenges in extending the original framework to more realistic problems including: accounting for slewing limitations, accounting for varying utility models affecting observation of the same point multiple times, and accounting for utility models preferring or not preferring proximal observations. We then describe realistic timing constraints required by DT. We also outline efforts to utilize deep learning to learn observation policies. In this paper we also describe efforts to flight validate DT technology on the ESA OPS SAT satellite. We also describe efforts to build an additional use case to study Planetary Boundary Layer (PBL) phenomena.

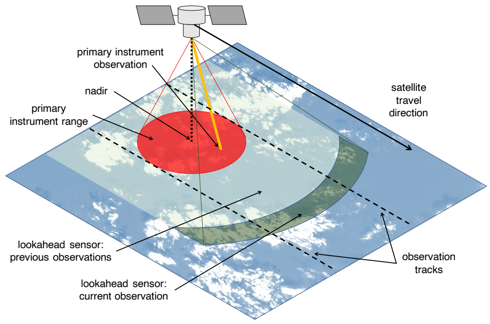

Figure 1. Dynamic targeting uses current and previous information from a lookahead sensor to identify future targets for the primary pointable instrument in order to improve science yield.

2. DYNAMIC TARGETING

Dynamic Targeting (DT) is a mission concept [1] (See Figure 1) in which a lookahead sensor acquires data which is then rapidly processed and used to retarget a sensor either to measure a designated science phenomena of interest (e.g. deep convective storm) or avoid conditions that may prevent measurement (e.g. clouds that obscure imaging of the Earth's surface or atmospheric measurements using lidar or IR sounder).

DT has been studied with multiple simulated datasets [1] [2] [3] [4] demonstrating great promise in improving science return for multiple use cases. DT algorithms for both lookahead data classification and targeting computation have even been tested onboard Snapdragon 855 embedded processors in the International Space Station (ISS) [5] [6] but not attached to any instruments (e.g. only batch processing on canned data - not on continuous data streams with processing time constraints). DT has even been used operationally for the JAXA TANSO-FTS-2 instrument flying on GOSAT-2 [13] for cloud avoidance.

Because the DT operations flow is a radical departure from traditional "mow the lawn" NASA Earth Science missions (see Table 1), it is critical that this new, complex operations flow be flight validated to provide assurance for future missions that DT is mature.

Because of the complexity of steps 1.-6. in Table 1 which must be performed continuously and within single overflight time constraints (typically 60 seconds or less for a typical Low Earth Orbit (LEO) geometry including slewing time), in-space flight demonstrations are needed to prove the technical feasibility of the DT concept.

2.1. Extending Slewing and Utility Models

Recent work on DT has addressed making the slew and utility models for DT more realistic. Initial DT work emerged from the SMICES mission concept which used a targeting radar for the primary instrument, and that the radar electronically steered and therefore had infinite agility. More recent work extends the DT models [7].

First, slewing constraints are enforced by a hard constraint on reachability by the next timestep. Longer slews are possible but are modelled by skipping an observation at the next timestep.

Second, the base DT model presumes that an observation can be acquired multiple times with constant utility unaffected by viewing geometry. This drives the base DT algorithm to repeatedly observe high utility targets in order to maximize utility. An extension to this model allows for utility to be affected by viewing geometry (either to prefer nadir or off nadir observations). Extensions also allow for increasing or decreasing utility from repeated observation of the same target. Additional work is investigating increasing (or non-decreasing) utility but only if sufficiently different viewing geometry. All of this additional complexity in utility makes the observation utility search space much more complex and computationally more challenging to search which is a challenge given that DT relies on the ability to respond (including instrument repointing) within a single overflight (typically 60 seconds at Low Earth Orbit). However with the advent of edge processors such as Snapdragon, NVIDIA, Gumstix, and others being flown in space [5] [6] [8] DT still appears computationally feasible.

2.2. Timing considerations for DT

One of the driving challenges for DT is that it must continually acquire and process data as well as retask the primary sensor all while flying at approximately 7.5 km per second (low Earth orbital speed). Table 2 shows a rough operational timeline for both the dedicated lookahead sensor case and the single sensor shared lookahead and primary (nadir) case.

As shown for a moderately agile spacecraft configuration (across track slew in 30s, along track slew of up to 50 degrees in 40s) the DT timeline is feasible.

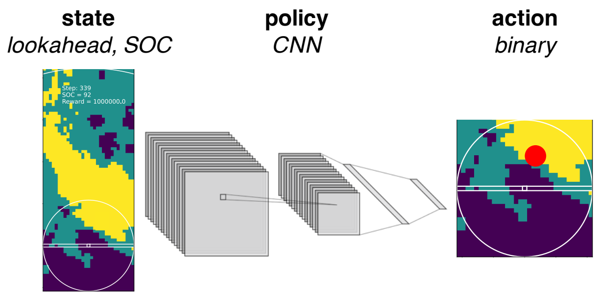

Figure 2. A convolutional neural network (CNN) is trained to learn policies using expert demonstrations. In this case, the CNN learns when to trigger observations as a function of the lookahead sensor data and the SOC.

3. LEARNING DT POLICIES

We are also investigating potentially improving DT algorithms [1] by leveraging learning and data-driven methods. Data patterns found in simulations can be exploited, allowing for custom policies that better adapt to different situations and scenarios (e.g., cloud avoidance vs. storm hunting). Specifically, we build on our DT algorithms so they can learn *when* it is best to trigger observations, while retaining prior logic that decides *where* to target the primary instrument. Herein we explore two different learning approaches: imitation learning and reinforcement learning.

Imitation learning consists of learning to replicate the actions of an expert agent. In our previous work, we present a dynamic programming (DP) algorithm that produces optimal policies (given a particular set of assumptions). However, DP is generally not deployable on missions because it is computationally expensive, it uses a lookahead sensor range that is physically unrealistic, and it requires information of future states. Despite these limitations, DP is very valuable as it not only serves as an upper bound on performance, but also as an expert whose optimal policies can be approximated offline via learning, allowing for onboard inference, planning, and execution with more realistic mission resources. We employ behavioural cloning, a form of imitation learning that predicts optimal actions from states using supervised learning on expert demonstrations, in this case from our DP algorithm. We train a convolutional neural network (CNN) where the input states are lookahead sensor data in conjunction with the current state of charge (SOC), while output actions are binary variables that dictate whether or not to trigger an observation (Figure 2).

Table 1. Traditional "mow the lawn" mission flow versus Dynamic Targeting (DT) operations flow. DNE = Does Not Exist

Element	"Mow the lawn" business as usual	Dynamic Targeting disruptive operations flow
1. Lookahead instrument and Data	DNE (does not exist)	Immediately streamed for real-time decision making
2. Edge Computing	Minimal if any	Required for real-time processing
3. Onboard processing	DNE - calibrated and corrected on ground	Lookahead data calibrated and corrected "on the fly" as needed
4. Lookahead data interpretation	DNE – analyzed on ground	Immediately analyzed into meaningful science categories for targeted instrument operations
5. Retargeting and reconfiguration of primary instrument	DNE – no reconfiguration and retargeting takes place	Decision-making on Reconfiguration and retargeting of targeted instrument
6. Targeted instrument	Stored onboard in raw format for later downlink	Reconfigured and targeted as per policy determined above; data stored onboard for later downlink.

*Table 2. Onboard timeline for processing of data and retargeting. * - tasks 3 and 4 can be done in parallel.*

Element	DT with dedicated lookahead sensor	DT utilizing primary sensor as lookahead sensor
1. Lookahead instrument acquires Data	4s	4s
2. Lookahead data transfer, calibration and analysis	4s	4s
3. Slew across track to detection*	30s	30s
4. Slew along track to nadir (1 instrument)*	N/A	40s
5. Primary instrument data take	4s	4s
Total of above timeline	42s	52s
Orbital Time for 40, 45, and 50 degree lookahead at 500km	60s, 74s, and 90s	60s, 74s, and 90s

Reinforcement learning is a different approach that learns policies without relying on an expert agent. Instead, a learning agent interacts with the environment and receives feedback via a reward function that incentivizes "good" actions and penalizes "bad" actions, in this case, by observing interesting or unimportant cloud types. Reinforcement learning is an iterative approach that improves gradually by learning from its mistakes and successes throughout many simulations runs. Here we use the Q-learning algorithm to train an agent so it learns when to trigger observations as represented by a binary action variable, while prior logic takes care of pointing the primary instrument.

4. FLIGHT OF DT ON OPS-SAT

Another version of DT has been developed that computes rewards over sets of targets while optimizing a cost function such as the accumulated slew angle.

The operational workflow of the so-called MIRE (dynaMIc taRgeting Experiment) version involves capturing a lookahead image, processing it to identify cloud formations, and selecting a subset of these clouds that maximizes the scientific reward. For each target, an image is captured. All these steps are performed within each 14-second DT cycle [9].

The European Space Agency's OPS-SAT cubesat [10] illustrated in Figure 3 is planned to be the first flight of

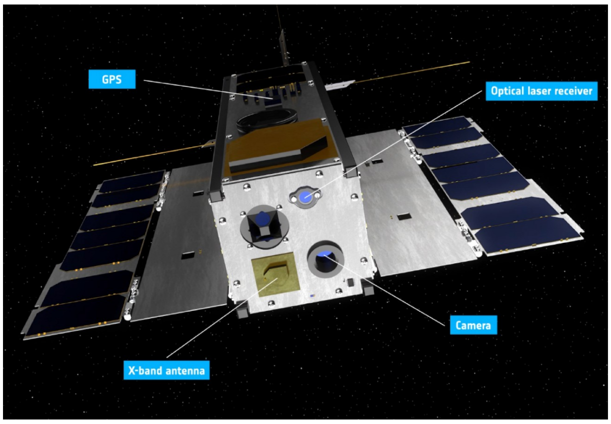

Figure 3. OPS-SAT Payload. Image courtesy of ESA.

DT by JPL. Unfortunately, the OPS-SAT does not have a lookahead sensor and the OPS-SAT pointing interface and agility do not allow for the single sensor to lookahead and slew back for the primary measurement. Therefore, DT on OPS-SAT utilizes the large field of view of the OPS-SAT imager and subframing to approximate the lookahead and primary imager. OPS-SAT acquires an image and subframes the ahead-of-track portion as the lookahead image and based on the contents of said image acquires a subsequent image subframing the behind-track portion as a stand-in for the primary image.

Following a successful trial run on the OPS-SAT Flatsat ground testbed in May 2023, the experiment is slated for on-board execution in September 2023. To further validate the algorithm's reliability over extended periods (multiple orbits), a second test has been scheduled before the end of the year.

5. DT FOR PLANETARY BOUNDARY LAYER STUDY

The DT team has also been studying the applicability of DT to capturing transient Planetary Boundary Layer (PBL) [11] [12] [13] phenomena. In this concept, the DT instrumentation may be distributed across multiple space platforms. The lookahead sensor(s) would include a Hyperspectral Infrared Sounder (Hyperspectral IR) and a Hyperspectral Microwave sounder (Hyperspectral MW) which would be used to search for areas of high variation in temperature and humidity that are indicative of PBL phenomena. This information would be used to target a DIfferential Absorption Lidar (DIAL) and Differential Absorption Radar (DAR). The concept is well suited to DT as the DIAL is narrow FOV and is best suited for clear sky applications and the Differential Absorption Radar (DAR) is also narrow FOV, is best suited for opaque storms, and is power hungry. GNSS data would also be used to supplement the PBL study.

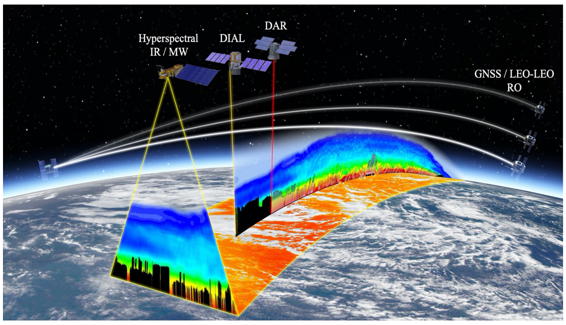

Figure 4. Concept for use of Dynamic Targeting to capture transient Planetary Boundary Layer phenomena. Image courtesy PBL Study Team.

6. RELATED WORK

Onboard analysis of science images has been flown before, even operationally. The Autonomous Sciencecraft on EO-1 analyzed Hyperion VSWIR images onboard for Thermal (Volcano, Wildfire) [14], Cryosphere [15], Cloud-screening, and Hydrology [16] applications as well as hyperspectral unmixing [17]. Notably the Cryosphere application used a Support Vector Machine (SVM) learned classifier. Later in the mission [18] Random decision forest (RDF) classifiers and salience were flown. However, ASE responses were on the order of one orbit (e.g. 90 minutes) and were not 60 seconds within the same overflight as proposed by DT. Notably, the Superpixel segmentation and Hyperspectral unmixing took 6 hours onboard the 6 MIPS Mongoose V processor. The IPEX mission demonstrated high throughput onboard data processing [8] (including RDF classifiers) using both an ATMEL ARM COTS processor as well as a Gumstix payload processor. But IPEX did not retarget or change any actions due to the data analysis and did not have a science quality instrument. ESA's recent Phi-Sat mission flew an Intel Myriad and used a deep learned CNN to detect clouds [19]. Most notably the TANSO-FTS-2 instrument on GOSAT-2 is operationally using cloud avoidance (a form of DT) [20].

Flight of onboard planners/schedulers onboard spacecraft is exceptionally rare. In 1999, the Remote Agent flew the RAX-PS planner which controlled the Deep Space One mission for two periods totalling approximately 24 hours [21]. In 2013, the CASPER planner flew onboard the IPEX cubesat [8] for over 1 year. More recently, the Mexec planner (related to the M2020 OBP) flew onboard the ASTERIA cubesat for 4-20 September 2019 [22]. Finally, the CASPER planner flew onboard the Earth Observing One mission from 2004-2017, controlling all EO-1 activities for over a dozen years [23].

7. CONCLUSIONS

Dynamic Targeting (DT) is a novel mission concept where lookahead sensor data is used continuously to target a primary sensor. DT has shown promise in simulation studies for storm hunting and cloud avoidance. In this paper we have described efforts to mature DT including extending the slewing and utility models, developing realistic timing requirements for Low Earth Orbit, learning DT policies, developing DT use cases for studying Palnetary Boundary Layer phenomena, and flight on ESA's OPS SAT orbital testbed.

ACKNOWLEDGMENTS

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

This work was sponsored by the NASA Earth Science Technology Office.

REFERENCES

[1] Alberto Candela, Jason Swope, and Steve Chien. Dynamic targeting to improve earth science missions. *Journal of Aerospace Information Systems (JAIS)*, 2023 (to appear).

[2] X. Bosch-Lluis, S. Chien, Q.Yue, J. Swope, P. Tavallali P., M. Ogut, I. Ramos, P. Kangaslahti, W. Deal, and C. Cooke. Developing radiometer and radar synergies using machine learning. In *International Geoscience and Remote Sensing Symposium*, July 2021.

[3] Mehmet Ogut, Xavier Bosch-Lluis, Pekka Kangaslahti, Isaac Ramos-Perez, Joan Francesc Munoz-Martin, Joelle Cooperrider, Qing Yue, Jason Swope, Peyman Tavallali, Steve Chien, Omkar Pradhan, William Deal, and Caitlyn Cooke. Autonomous capabilities and command and data handling design for the smart remote sensing of cloud ice. In *International Geoscience and Remote Sensing Symposium (IGARSS 2022)*, Kuala Lumpur, Malaysia, July 2022.

[4] Zaki Hasnain, James Mason, Jason Swope, Joshua Joshua Vander Hook, and Steve Chien. Agile spacecraft imaging algorithm comparison for earth science. In *International Workshop on Planning & Scheduling for Space (IWPSS)*, July 2021. URL `https://ai.jpl.nasa.gov/public/papers/Hasnain-IWPSS2021-paper-13.pdf`.

[5] Emily Dunkel, Jason Swope, ALberto Candela Garza, Lauren West, Steve Chien, Zaid Towfic, Leonie Buckley, Juan Romero-Canas, Jose Luis Espinosa-Aranda, Elena Hervas-Martin, and Mark Fernandez. Benchmarking deep learning models on myriad and snapdragon processors onboard the iss. *Journal of Aerospace Information Systems*, 140 (2021) 103758, 2023 (to appear). URL `https://doi.org/10.1016/j.robot.2021.103758`.

[6] Jason Swope, Faiz Mirza, Emily Dunkel, Alberto Candela, Steve Chien, Alexandra Holloway, Damon Russell, Joe Sauvageau, Douglas Sheldon, and Mark Fernandez. Benchmarking space mission applications on the snapdragon processor onboard the iss. *Journal of Aerospace Information Systems*, 2023 (to appear).

[7] Akseli Kangaslahti, Alberto Candela, Jason Swope, Qing Yue, and Steve Chien. Dynamic targeting of satellite observations incorporating slewing costs and complex observation utility. In *In review*, 2023.

[8] S. Chien, J. Doubleday, D. R. Thompson, K. Wagstaff, J. Bellardo, C. Francis, E. Baumgarten, A. Williams, E. Yee, E. Stanton, and J. Piug-Suari. Onboard autonomy on the intelligent payload experiment (ipex) cubesat mission. *Journal of Aerospace Information Systems (JAIS)*, April 2016.

[9] Juan M. Delfa Victoria, Alberto Candela, and Steve A. Chien. Enhanced dynamic targeting for the opssat cubesat. In *Proc. of International Workshop on Planning and Scheduling for Space*, July 2023.

[10] ESA. OPS-SAT web. `https://www.esa.int/Enabling_Support/Operations/OPS-SAT`, accessed on 2023-03.

[11] J. Teixeira, J. R. Piepmeier, A. R. Nehrir, C. O. Ao, S. S. Chen, C. A. Clayson, A. M. Fridlind, M. Lebsock, W. McCarty, H. Salmun, J. A. Santanello, D. Turner, Z. Wang, and X. Zeng. *Toward a Global Planetary Boundary Layer Observing System: The NASA PBL Incubation Study Team Report*. NASA PBL Incubation Study Team, 2021.

[12] Planetary Boundary Layer. `https://www.sciencedirect.com/topics/earth-and-planetary-sciences/planetary-boundary-layer`, 2023.

[13] Planetary Boundary Layer Targeted Observable. `https://science.nasa.gov/earth-science/decadal-pbl`, 2023.

[14] A. Davies, S. Chien, V. Baker, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. Castano, B. Cichy, R. Lee, G. Rabideau, D. Tran, and R. Sherwood. Monitoring active volcanism with the autonomous sciencecraft experiment (ase). *Remote Sensing of Environment*, 101 (4):427–446, April 2006.

[15] T. Doggett, R. Greeley, A. Davies, S. Chien, B. Cichy, R. Castano, K. Williams, V. Baker, J. Dohm,

and F. Ip. Autonomous on-board detection of cryospheric change. *Remote Sensing of Environment*, 101 (4):447–462, 2005.

[16] F. Ip, J. Dohm, V. Baker, T. Doggett, A. Davies, R. Castano, S. Chien, B. Cichy, R. Greeley, and R. Sherwood. Development and testing of the autonomous spacecraft experiment (ase) floodwater classifiers: Real-time smart reconnaissance of transient flooding. *Remote Sensing of Environment*, 101 (4):pp. 463–481., 2005.

[17] D. R. Thompson, B. Bornstein, S. Chien, S. Schaffer, D. Tran, B. Bue, R. Castano, D. Gleeson, and A. Noell. Autonomous spectral discovery and mapping onboard the eo-1 spacecraft. *IEEE Transactions on Geoscience and Remote Sensing*, 2012.

[18] K. Wagstaff, S. Chien, A. Altinok, U. Rebbapragada, D. Thompson, S. Schaffer, and D. Tran. 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, July 2018. URL `https://ai.jpl.nasa.gov/public/papers/wagstaff-ijcai2017-novelty.pdf`. Also appears at AI in the Oceans and Space Workshop, International Joint Conference on Artificial Intelligence (IJCAI 2017).

[19] Gianluca Giuffrida, Luca Fanucci, Gabriele Meoni, Matej Batič, Léonie Buckley, Aubrey Dunne, Chris van Dijk, Marco Esposito, John Hefele, Nathan Vercruyssen, et al. The ϕ-sat-1 mission: The first onboard deep neural network demonstrator for satellite earth observation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2021.

[20] Hiroshi Suto, Fumie Kataoka, Nobuhiro Kikuchi, Robert O Knuteson, Andre Butz, Markus Haun, Henry Buijs, Kei Shiomi, Hiroko Imai, and Akihiko Kuze. 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, 2021.

[21] Ari K Jónsson, Paul H Morris, Nicola Muscettola, Kanna Rajan, and Benjamin D Smith. Planning in interplanetary space: Theory and practice. In *Proceeding of Conference on Artificial Intelligence Planning Scheduling*, pages 177–186, 2000.

[22] M. Troesch, F. Mirza, K. Hughes, A. Rothstein-Dowden, R. Bocchino, A. Donner, M. Feather, B. Smith, L. Fesq, B. Barker, and B. Campuzano. Mexec: An onboard integrated planning and execution approach for spacecraft commanding. In *Workshop on Integrated Execution (IntEx) / Goal Reasoning (GR), International Conference on Automated Planning and Scheduling (ICAPS IntEx/GP 2020)*, October 2020. URL `https://ai.jpl.nasa.gov/public/papers/IntEx-2020-MEXEC.pdf`.

[23] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, and D. Boyer. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication (JACIC)*, pages 196–216, April 2005. URL `https://ai.jpl.nasa.gov/public/papers/chien-JACIC2005-UsingAutonomy.pdf`.