

# Artificial Intelligence-Based Challenges as an Educational Tool in Aerospace Engineering: the u3S Laboratory Experience

A. Lotti\*, D. Modenini\*\*

\**Department of Industrial Engineering, Alma Mater Studiorum Università di Bologna  
Via Fontanelle 40 47121 Forlì Italy (alessandro.lotti4@unibo.it).*

\*\**Department of Industrial Engineering, Interdepartmental Centre for Industrial Research in Aerospace  
Alma Mater Studiorum Università di Bologna  
Via Fontanelle 40 47121 Forlì Italy (e-mail: dario.modenini@unibo.it)*

**Abstract:** This manuscript reports on the positive impact that Artificial Intelligence based challenges are having on the educational and research activities of the u3S laboratory at the University of Bologna, within the framework of degree courses in aerospace engineering. Three examples are discussed from the past academic years, highlighting the benefits along with the challenges, of using AI based competitions as the source of hand-on education to carry out final projects.

Copyright © 2024 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

**Keywords:** CubeSats, satellite testing, hands-on education, laboratory equipment.

## 1. INTRODUCTION

Artificial Intelligence (AI) is gaining impressive momentum in the aerospace engineering field. Applications are the most diverse, encompassing, design, optimization, systems autonomy, operations, data analysis, air and space traffic management, just to mention a few. If on one side, the research and technological applications are being comprehensively reported in the literature, the role of AI as a tool for aerospace education has been receiving less attention.

An emerging trend that can be observed is the publication of AI-based challenges. Indeed, open challenges are an effective way to boost research, as they inherently promote collaboration and problem solving across disciplines.

In this manuscript, we describe the experience gained by the Microsatellite and Space Microsystems (u3S) Laboratory at the University of Bologna with a number of AI-based challenges in the aerospace domain fields. Participation to those challenges represent a fruitful source for students' internships or final projects, as well as a self-training and learning tool for PhD students at the beginning of their research path.

## 2. SPACE ENGINEERING EDUCATIONAL ACTIVITIES AT THE u3S LAB

Since 2003, the Alma Mater u3S lab at the University of Bologna has been conducting research in space technology. The laboratory hosts numerous students from the Aerospace Engineering courses, aiming to train them in space systems design through various hands-on projects. Focusing on small satellite platforms, the lab aligns with university budget constraints and educational goals, offering a cost-effective test-bed for students to learn by doing. This approach allows students to engage in the entire space project lifecycle, including on-orbit operations.

The lab's activities began with the ALMASat-1 project (Tortora et al., 2006), a microsatellite launched in 2012 on board the VEGA maiden flight, involving over 60 BSc, MSc, and PhD students in its development. Building on the success of ALMASat-1, the laboratory later supported the development and operations of the European Student Earth Orbiter (ESEO) microsatellite (Curzi et al., 2020), an ESA Education Office's led mission launched in 2018, having tens of undergraduate, graduate, and PhD students involved.

More recently, the focus of the lab moved towards the following three target areas in space systems engineering:

1. spacecraft navigation, either GNSS-based or imaged-based (Locarini et al., 2016).
2. attitude determination and control (Bahu et al., 2020; Modenini et al., 2021);
3. space traffic management and collision avoidance (Modenini et al., 2022; Curzi et al., 2023).

Starting from 2020, the lab began gaining experience on the use of AI in aerospace engineering problems pertaining to the first two research area above, including:

- deep-learning for real-time pose estimation in the context of On-Orbit Servicing and Active Debris Removal missions (Lotti et al., 2023a);
- enhancing orbit predictions from two-line-element sets using neural networks (Curzi et al., 2022).

The first application, in particular, fostered the begin of u3S lab involvement into AI-based challenges.

In the remaining of the manuscript, we will present the outcome of four challenges we participated in the past Academic Years, namely:

- AY 2020-21 and AY 2021/2022: Pose estimation challenges.
- AY 2021-22: HYPERVIEW challenge.
- AY 2022-23: OrbitalAI challenge.

### 3. THE SATELLITE POSE ESTIMATION COMPETITIONS

The Satellite Pose Estimation Competitions were organized by the Space Rendezvous Laboratory (SLAB), Stanford University, and the Advanced Concepts Team at the European Space Agency (ESA) in 2019 and 2021 (Kisantal et al., 2020; Park et al., 2023). The purpose of these competitions was to raise awareness within the research community about the need to develop algorithms for accurate pose estimation of a potentially uncooperative target from monocular images captured by an active chaser. While pose estimation is extensively researched by the computer science community, space presents additional challenges, such as harsh illumination conditions, low image quality, and dataset scarcity, which require specific solutions.

In this framework, the pose is defined as the quaternion and position vector that align the target's body frame with the chaser's camera frame, see Fig. 1. The pose error, i.e., the difference between the ground truth and the estimated pose, can be computed as:

$$E = \frac{1}{N} \sum_{i=1}^N \left( \frac{|\mathbf{t}_i - \hat{\mathbf{t}}_i|_2}{|\mathbf{t}_i|_2} + 2 \arccos(|\langle \mathbf{q}_i, \hat{\mathbf{q}}_i \rangle|) \right) \quad (1)$$

Where  $\mathbf{t}_i$ ,  $\mathbf{q}_i$  and  $\hat{\mathbf{t}}_i$ ,  $\hat{\mathbf{q}}_i$  are the ground truth and estimated position vectors and attitude quaternions on the  $i$ -th image, while  $N$  is the number of images.

#### 3.1 First edition

The first competition leveraged the Spacecraft Pose Estimation Dataset (SPEED) (Sharma et al. 2019), which features 15000 computer-generated images of the Tango spacecraft at distances ranging from 3 m to 50 m. Our laboratory began research in this field after the challenge had concluded, allowing us to benefit from the dataset and competition results to gain experience on the topic. In this framework, a Ph.D. student at the beginning of their 3-year research program, developed a pose estimation pipeline (Lotti et al., 2023a). Although the structure of this pipeline was not novel, it specifically investigated the accuracy achievable on edge devices. This aspect was overlooked in the challenge, where the evaluation metric focused solely on accuracy without considering computational burden.

The algorithm was based on a three-step processing: 1. Object detection, 2. Landmark regression, and 3. Pose estimation. A fundamental assumption is the knowledge of the 3D model of the target. Although this model was not provided, it can be reconstructed by manually selecting a coherent set of distinctive features in the images and leveraging the known ground truth poses. This assumption allows to reduce the pose

estimation problem to a landmark regression problem. In principle, knowing the 3D keypoints, their location in the image and the camera intrinsics is sufficient to estimate the relative pose. Two distinct neural networks were employed for the first two steps, trained in supervised manner. The 3D model, together with the poses, was used to automate the generation of the training labels, consisting of bounding box and landmarks coordinates. For the first step, an off-the-shelf MobileDet (Xiong et al., 2021) detector was employed. The second step involved a backbone network equipped with a regression head trained to directly regress normalized landmarks coordinates. EfficientNet Lite (Tan et. al, 2019) Convolutional Neural Networks (CNNs) were selected as the backbone network, allowing for the creation of a family of regression models tailored to process input crops of increasing size, thus retaining more information. Pose estimation was then handled using the well-known EPnP algorithm (Lepetit et. al, 2009).

The following optimization have then been applied to the neural networks: 1. Conversion to TensorFlow-Lite, 2. Quantization, and 3. Co-compilation for the edge Tensor Processing Unit (TPU). Since the test labels of SPEED were not released, the pipeline was evaluated against the training set using a 4-fold cross-validation setup. The lightest version of the algorithm, featuring a total of 6.9 million parameters across the two networks, achieved an average pose error of about  $0.8^\circ$  and 5.3 cm, scoring  $E = 0.0186$ . In comparison, the competition winner scored  $E = 0.0094$  (evaluated on the competition test set) but featured about 25 times more network parameters. The lightest model processed 7.7 frames per second on the TPU with a measured power consumption of about 2.2 W. The heaviest model achieved an improvement of about 23% in the competition pose error but processed only about 2 fps. This speed drop is attributed to the fact that the larger landmark regression network does not fit entirely into

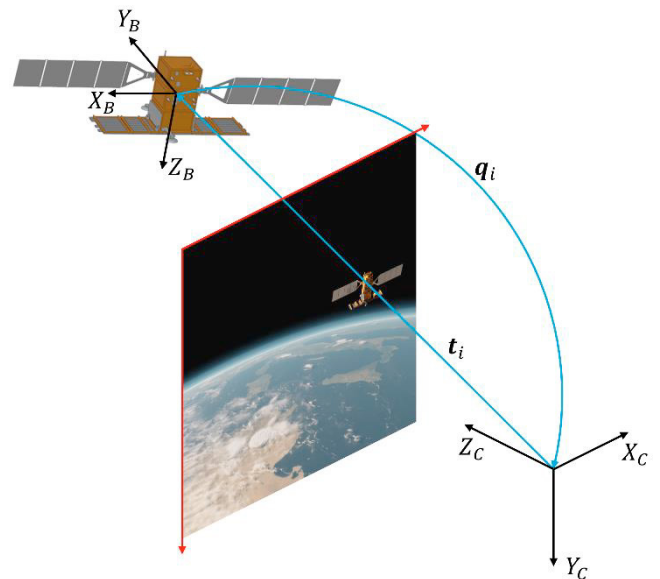


Fig. 1. Exemplification of the monocular pose estimation process

the edge TPU chip memory, highlighting the importance of software-hardware co-design.

### 3.2 Second edition

A second edition of the challenge was held in 2021. The problem statement and the evaluation metric remained the same, but the SPEED+ dataset (Park et al., 2022) was adopted, adding the domain gap obstacle to the competition. Domain gap refers to the decrease in accuracy that deep neural networks experience when they are applied to a different but related dataset. Collecting a large amount of accurately labeled spacecraft images in orbit is both technologically challenging and economically unfeasible. Therefore, the common approach is to train neural networks on synthetic images. Despite advances in render engines with physically based ray-tracing models, accurately modeling all possible conditions a camera might face in orbit remains challenging. Additionally, the target may have a different appearance due to the extended time spent in orbit. In other words, while the algorithm might perform well on synthetic data, an accuracy drop may occur on real images because the network has developed an over-reliance on synthetic features.

To address this issue, in the second edition, participants were asked to train their algorithms on rendered images of the Tango spacecraft and test them on two sets of real pictures captured in a laboratory environment using a satellite mockup. These sets aimed to faithfully emulate Earth's albedo (lightbox category) and extreme lighting conditions (sunlamp category), which are challenging to accurately replicate in rendering software.

The same Ph.D. student proposed an improved version of the three-stage pipeline (Lotti et al., 2023b), specifically tailored for the domain gap. In particular, both CNNs were replaced by Vision Transformer architectures, specifically SWIN Transformers (Liu et al., 2021). This choice was driven by the improved generalization ability of these models on many downstream tasks. Additionally, an adversarial training approach was adopted. Both the detection and regression networks were equipped with a classifier. During each training step, a single batch contained exactly 50% real and 50% synthetic images. The networks were tasked with regressing bounding box and landmark coordinates and simultaneously classifying the images as either real or synthetic using binary cross-entropy (i.e., no distinction was made between sunlamp and lightbox categories) through a dedicated classifier subnetwork.

Following the traditional adversarial learning scheme, the opposite of the discriminator loss was propagated through the backbone. This forced the backbone to elaborate features that are non-discriminative based on the domain. The learning objectives are summarized in (2)-(4), where  $W$  stands for network weights,  $L$  for the loss,  $CLS$  and  $REG$  for classification and regression respectively, and  $\lambda$  is a scaling parameter. R, C, and B stand for regressor, classifier, and backbone networks respectively.

$$\widehat{W}_C = \text{minimize}_{W_C}(L_{CLS}) \quad (2)$$

$$\widehat{W}_R = \text{minimize}_{W_R}(L_{REG}) \quad (3)$$

$$\widehat{W}_B = \text{minimize}_{W_B}(L_{REG} - \lambda L_{CLS}) \quad (4)$$

Since real image poses were not supplied, the regression supervision was provided only on the synthetic data.

Throughout the competition, different model sizes and data augmentations were tested. The best model, featuring SWIN base models for both detection and landmark regression, achieved 5th place on the lightbox leaderboard and 4th place on the sunlamp leaderboard, scoring  $E = 0.224$  and  $E = 0.141$  respectively. These results, compared to the ones obtained in the previous competition, highlight the challenge posed by the domain gap.

### 3. THE HYPERVIEW CHALLENGE

The Hyperview Challenge (Nalepa et al., 2022) is a joint initiative by ESA  $\Phi$ -Lab, KP Labs, and QZ Solutions aimed at promoting the development of algorithms for automating the estimation of soil parameters from hyperspectral space images to support more sustainable agriculture. Knowledge of certain terrain characteristics is crucial for selecting the correct mix of fertilizers, which can also help reduce their overall use, benefiting both farmers' costs and the environment. Currently, this process is performed manually: soil samples are collected in the field and sent to specialized laboratories for chemical analysis. This method is costly, time-consuming, and not scalable in terms of both space and time, providing farmers with a limited and infrequently updated picture of their fields' condition.

The challenge asked participants to regress soil characteristics from airborne hyperspectral images collected in 2021 over Poland. The camera featured 150 continuous spectral bands, ranging from visible to near-infrared, reflecting the characteristics of the Intuition 1 satellite's payload, where the winning solution would have been run. Specifically, participants were tasked with regressing the soil's pH and the concentrations of potassium (K), phosphorus pentoxide ( $P_2O_5$ ), and magnesium (Mg). Ground truth parameters were obtained through in-situ sampling and laboratory measurement. The competition dataset was quite challenging, comprising only 1732 training images with varying dimensions, from 11 px to 284 px. Submissions were evaluated based on the average ratio of the Mean Squared Error (MSE) for each soil parameter prediction to the MSE of a baseline model. The baseline model simply predicted the average parameter value from the training set.

The competition ran from February to July 2022. The u3S lab participated with a team composed of a Bachelor's student and a Ph.D. student, achieving 8th place on the final leaderboard, which improved to 4th place when considering the code reproducibility requirement. Specifically, the Bachelor's student had the opportunity to conduct both an internship and its final thesis on the competition, as well as further develop the algorithms after the competition ended, in collaboration with and under the supervision of the Ph.D. candidate.

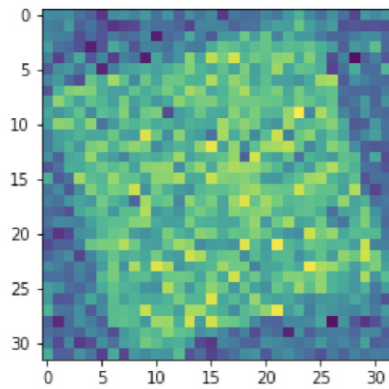


Fig. 2. Sample band patch corrupted by noise (Ballabeni 2022)

Given that AI topics are not taught within the Bachelor's program, the student began the internship by following online courses on the relevant topics and conducting a literature review on the proposed problem. Subsequently, Python codes were developed to analyze the dataset, followed by preprocessing functions, training and inference notebooks. In the final version of the algorithm, it was decided to retain all the hyperspectral channels while constraining the spatial resolution. The input size of the neural network was set to 32x32x150 px. For patch sizes smaller than 16 px, the patch was tiled and eventually flipped, while bilinear resizing was employed for larger patches to fit the input shape. Gaussian noise was also added, see Fig. 2. While this could potentially corrupt the information, it proved beneficial in overcoming network overfitting. The model developed leverages an EfficientNet-lite0 backbone, chosen to accommodate the computational constraints associated with onboard inference.

Cross-validation was employed to validate the model on the train set. The results confirmed that using small patches balances accuracy and latency, despite outliers linked to infrequently represented soil conditions. The algorithm, tested on a Coral Dev Board Mini CPU, processed 180 frames per second after TensorFlow-Lite optimization, proving its efficiency for on-orbit processing.

## 5. THE ORBITAL AI CHALLENGE

Given the success of the previous experience, in 2023 the u3S lab also participated in the OrbitalAI Challenge<sup>1</sup>, providing an internship and bachelor's thesis opportunity. The competition was set up again by ESA Φ-Lab, in partnership with Thales Alenia Space and Microsoft. This competition differed from traditional AI-based competitions in that submissions were not scored according to a numerical metric. Instead, participants had to provide an application leveraging onboard AI and visible/hyperspectral sensors to be eventually run on the ISS Mounted Accessible Global Imaging Node (IMAGIN-e), an edge compute testbed equipped with visible and hyperspectral sensors for validating remote-sensing applications. Final submissions were scored by a commission based on the following criteria: 1. Suitability of the solutions for solving

societal and/or business challenges, 2. Creativity, 3. Technical feasibility, 4. Contribution to open science.

The competition was organized into four milestones: 1. Team profile, 2. Idea concept, 3. Initial model description, 4. Final submission. The first three milestones were not compulsory for the final evaluation but proved helpful in ensuring the project remained on track. Participants were encouraged to use simulated images obtained from Sentinel and PRISMA missions, reprocessed through a Python notebook provided by the organizers, aimed at adapting images as if they were captured from the IMAGIN-e camera. The final submission included a working prototype, a pitch, and a presentation.

The u3S lab team proposed HyperLSTM, an algorithm for short-term predictions about the evolution of chlorophyll-A concentration in water bodies, which could be exploited to quickly spot and monitor potential hazards to the environment before they happen. Indeed, Chlorophyll-A concentration is one of the principal indicators of water quality. Traditional in situ measuring is time-consuming and often too slow to predict the ever-changing ecosystem, as large algae blooms can appear from one day to the next.

During the internship and final thesis, a first prototype was developed. The idea was to exploit a CNN Long Short Term Memory (LSTM) neural network that, starting from two images of the same area captured approximately three days apart, in compliance with ISS revisit times, could predict its evolution three days ahead. During training, images were preprocessed by applying the NDCI - Normalized Difference Chlorophyll Index, which allows for reducing the information in the multispectral channels into a single scalar index employing two spectral bands, thus benefiting computational efficiency. The network was trained to regress this same index, using 256x256 px image maps, see Fig. 3. Throughout the internship, a segmentation model to distinguish between land and water areas was also developed.

Despite the short duration of the challenge and the fact that AI-related concepts were new to the student, he was able to

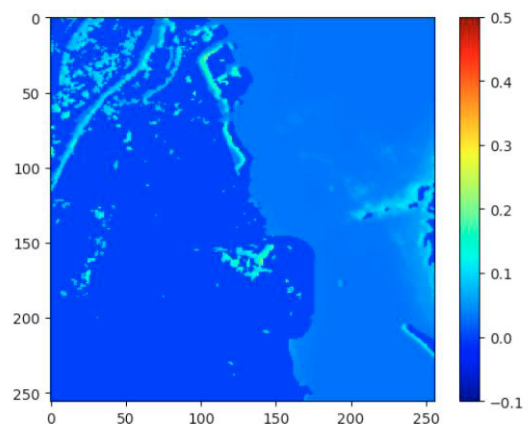


Fig. 3. Sample image map of the NDCI used to train the LSTM, (Gianfelici 2023)

<sup>1</sup> <https://ai4eo.eu/portfolio/orbitalai-challenge/>.

submit a solution within the deadline and the evaluation board deemed it satisfactory.

## 6. CONCLUSIONS

International challenges are usually established to advance the state of the art in specific fields. Our experience with AI-related challenges within a degree course in Aerospace engineering at the University of Bologna has shown that they are also effective in engaging students in hands-on educational activities. The scarce, if any, background in AI related disciplines possessed by the involved students revealed to be, rather than a limiting factor, a strong motivation to broaden their knowledge and skills set.

This positive, yet somehow unexpected, outcome highlights the potential of such challenges as a tool to educate future aerospace engineering professionals, by driving students to acquire new competencies in AI which are increasingly relevant across multiple domains.

As a result, we plan to continue and possibly further increase our efforts to involve students in such proficient training activities during the next academic years.

## REFERENCES

- Ballabeni, A. (2022). Implementazione di reti neurali per missioni di telerilevamento satellitare da immagini iperspettrali. *University of Bologna - Forlì* <http://amslaurea.unibo.it/view/cds/CDS9234/>.
- Bahu, A., and Modenini, D. (2020). Automatic mass balancing system for a dynamic CubeSat attitude simulator: development and experimental validation. *CEAS Space Journal*, 12, 597-611.
- Curzi, G., Lucci, A., Locarini, A., Modenini, D., Tortora, P., Mariotti, G., Cinarelli, D., Fabbri, V., and Melega, N. (2020). Lessons Learnt From Operating ESEO Educational Spacecraft. In *Proceedings of the 3rd Symposium on Space Educational Activities*, 279-283.
- Curzi, G., Modenini, D., and Tortora, P. (2022). Two-line-element propagation improvement and uncertainty estimation using recurrent neural networks. *CEAS Space Journal*, 14(1), 197-204.
- Curzi, G., and Modenini, D. (2023). Analytic solution for perturbed Keplerian motion under small acceleration using averaging theory. *Advances in Space Research*, 72(7), 2445-2459.
- Gianfelici, R. (2023). Implementazione di reti neurali per la previsione di parametri di qualità dell'acqua da immagini satellitari. *University of Bologna - Forlì* <http://amslaurea.unibo.it/view/cds/CDS9234/>
- Kisantal, M., Sharma, S., Park, T. H., Izzo, D., Märten, M., and D'Amico, S. (2020). Satellite pose estimation challenge: Dataset, competition design, and results. *Transactions on Aerospace and Electronic Systems*, IEEE, 56(5), 4083-4098.
- Lepetit, V., Moreno-Noguer, F., and Fua, P. (2009). EPnP: An Accurate O(n) Solution to the PnP Problem. *International Journal of Computer Vision*, volume 81, No. 2, 155–166.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B. (2021). Swin transformer: hierarchical vision transformer using shifted windows. In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Montreal, QC, Canada, 9992–10002.
- Locarini, A., Tortora, P., & Avanzi, A. (2016). Development and test campaign of a GPS-based navigation sub-system for the ESEO mission. In *2016 Metrology for Aerospace (MetroAeroSpace)*, IEEE, 567-572.
- Lotti, A., Modenini, D., Tortora, P., Saponara, M., & Perino, M. A. (2023). Deep learning for real-time satellite pose estimation on tensor processing units. *Journal of Spacecraft and Rockets*, AIAA, 60(3), 1034-1038.
- Lotti, A., Modenini, D., Tortora, P. (2023). Investigating Vision Transformers for Bridging Domain Gap in Satellite Pose Estimation. In: Ieracitano, C., Mammone, N., Di Clemente, M., Mahmud, M., Furfaro, R., and Morabito, F.C. (eds.), *The Use of Artificial Intelligence for Space Applications*, 299–314. Springer Nature Switzerland, Cham.
- Modenini, D., Bahu, A., and Tortora, P. (2021). Nanosatellite-class dynamic attitude simulator for hands-on aerospace control education. *IFAC-PapersOnLine*, 54(12), 38-42.
- Modenini, D., Curzi, G., Locarini, A. (2022). Relations between collision probability, Mahalanobis distance, and confidence intervals for conjunction assessment. *Journal of Spacecraft and Rockets*, 59(4), 1125-1134.
- Nalepa, J., Le Saux, B., Longepe, N., Tulczyjew, L., Myller, M., Kawulok, M., Smykala, K., Gumiel, M. (2022). The Hyperview Challenge: Estimating Soil Parameters from Hyperspectral Images. In *2022 IEEE International Conference on Image Processing (ICIP)* 4268–4272. IEEE, Bordeaux, France.
- Park, T.H., Märten, M., Lecuyer, G., Izzo, D., D'Amico, S. (2022). SPEED+: Next-Generation Dataset for Spacecraft Pose Estimation across Domain Gap. In: *2022 IEEE Aerospace Conference (AERO)*, 1–15.
- Park, T.H., Märten, M., Jawaid, M., Wang, Z., Chen, B., Chin, T.-J., Izzo, D., D'Amico, S. (2023). Satellite Pose Estimation Competition 2021: Results and Analyses. *Acta Astronautica*, 204, 640–665.
- Sharma, S., and D'Amico, S. (2019). Pose Estimation for Non-Cooperative Rendezvous Using Neural Networks. In *2019 AAS/AIAA Astrodynamics Specialist Conference*, American Astronautical Soc. Paper 19-350, Springfield, VA, 1–20.
- Tan, M., and Le, Q. (2019). EfficientNnet: Rethinking Model Scaling for Convolutional Neural Networks. In

*International Conference on Machine Learning*, Univ. of Cambridge, Cambridge, 6105–6114.

Tortora, P., Giulietti, F., Corbelli, A., and Fabbri, V. (2006). ALMASat Attitude Control Hardware-In-The-Loop Simulations. In *57th International Astronautical Congress* (pp. C1-P).

Xiong, Y., Liu, H., Gupta, S., Akin, B., Bender, G., Wang, Y., Kindermans, P.-J., Tan, M., Singh, V., Chen, B. (2021). MobileDets: Searching for Object Detection Architectures for Mobile Accelerators. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, New York, 3824–3833.