

# HYPERBOLA (HYPerspEctRal onBOard cLoud AI): Preliminary Results

M. Banti<sup>a</sup>, G. Gousios<sup>a</sup>, T. Mamouka<sup>a</sup>, P. Vourlioti<sup>a</sup>, S. Neophytides<sup>b</sup>, C. Paraskevas<sup>a</sup>, S. Kotsopoulos<sup>a</sup>, M. Mavrovouniotis<sup>b</sup>,

<sup>a</sup>Ahtri tn 12, Kesklinna linnaosa, Tallinn, Harju maakond, 15551, Estonia

<sup>b</sup>ERATOSTHENES Centre of Excellence, 82 Franklin Roosevelt, 3012, Limassol, Cyprus

[vvourlioti@neuralio.space](mailto:vvourlioti@neuralio.space)

## ABSTRACT

Cloud and cloud-shadow masking in hyperspectral (HS) remote sensing is a critical step for reliable Earth Observation data. However, performing this task on-board a satellite remains challenging. Traditional approaches, such as threshold-based tests or CNN classifiers, often fail to capture complex spectral-spatial dependencies and typically treat cloud detection as a binary problem, overlooking cloud-shadow and multi-class distinctions. We propose HYPERBOLA (HYPerspEctRal onBOard cLoud AI), an AI framework that integrates Vision Transformers (ViTs), spectral-spatial Graph Neural Networks (GNNs), and hybrid CNN-Transformer architectures to improve cloud and cloud-shadow segmentation. In this work, we present preliminary results focusing on a single hybrid CNN-Transformer model (Hybrid Model 2). This architecture shows promising performance in segmenting Sea and Land classes, while highlighting current challenges in detecting thin or sparse clouds. These findings establish a baseline for onboard hyperspectral cloud masking and demonstrate the feasibility of transformer-based methods in constrained environments. Future work will extend the evaluation to the full set of proposed models (ViTs, GNNs, and hybrid variants), incorporate advanced optimization techniques such as structured pruning and 8-bit quantization, and benchmark performance across a broader set of test scenes. The ultimate goal of HYPERBOLA is to deliver an efficient and accurate onboard AI framework for real-time hyperspectral cloud and cloud-shadow segmentation.

**Keywords:** Hyperspectral Imaging, Cloud Masking, On-board AI, Model Optimization

## 1. INTRODUCTION

HyperSpectral imaging (HSI) has transformed Earth Observation by enabling the acquisition of detailed spectral profiles for every pixel in a scene. With hundreds of narrow and contiguous spectral bands spanning the visible to shortwave infrared, hyperspectral sensors offer unparalleled capabilities for applications such as vegetation health monitoring, mineral exploration, maritime surveillance, and precision agriculture. This spectral richness allows for fine-grained material discrimination far beyond what traditional multispectral systems can achieve.

However, one of the most persistent challenges in the analysis of hyperspectral satellite data is the presence of clouds and cloud shadows, which significantly degrade the quality and utility of the captured imagery. Cloud-contaminated pixels obstruct surface visibility, reduce the effective data yield, and interfere with downstream analysis such as classification, anomaly detection, and temporal monitoring [1]. Cloud masking—the process of identifying and excluding such pixels—is therefore a crucial preprocessing step in hyperspectral pipelines.

As noted, cloud masking is critical for hyperspectral imagery, particularly when dealing with challenges such as thin clouds, cloud shadows, and complex terrain. Our approach leverages Vision Transformers (ViTs) to capture long-range spectral dependencies, effectively handling subtle cloud types like cirrus clouds. The Spectral-Spatial Graph Neural Networks (GNNs) are used to model pixel-level spatial relationships, allowing for more precise cloud-shadow segmentation. Meanwhile, hybrid CNN-Transformer architectures provide a balanced solution, combining local texture recognition with global context, which is essential for complex backgrounds such as sunglint on water or urban areas.

Conventional cloud detection techniques typically rely on thresholding rules applied to brightness, temperature, or spectral indices. While effective in certain conditions, these methods struggle with subtle cloud types (e.g., thin cirrus), complex backgrounds (e.g., sunglint on water), or challenging terrain (e.g., snow or urban areas), especially when applied to hyperspectral data that has a much higher dimensionality than traditional RGB or multispectral imagery [2]. In recent years, deep learning models—particularly Convolutional Neural Networks (CNNs)—have become the standard for cloud

detection and segmentation in satellite imagery [3]. These models are typically trained on RGB or multispectral inputs and are deployed on ground-based systems. However, as the size and frequency of Earth Observation missions increase, the burden of storing and transmitting full-resolution hyperspectral cubes from orbit becomes a limiting factor. This motivates a shift toward onboard AI-based processing, which allows satellites to analyze data in real-time and transmit only the most relevant outputs to ground stations [4][5].

Several studies have proposed lightweight deep learning models for on-orbit inference, including 1D-CNN architectures that process spectral vectors without full 2D spatial convolution [5]. Notably, the 1D-Justo-LiuNet model, which was deployed on the HYPSONO-1 satellite, demonstrated the potential of efficient onboard segmentation, achieving over 93% accuracy while maintaining a model size of just 4.5k parameters [6]. Such architectures outperform more complex 2D-CNNs and lightweight Vision Transformers in constrained environments, providing a strong baseline for onboard hyperspectral semantic segmentation [6].

While recent breakthroughs, such as the in-orbit deployment of the 1D-Justo-LiuNet model on the HYPSONO-1 satellite, have demonstrated the feasibility of real-time, onboard semantic segmentation of hyperspectral imagery into sea, land, and cloud classes [1], significant gaps remain in achieving broader operational autonomy. These include challenges in multi-class generalization (e.g., differentiating thin clouds, cloud shadows, or snow), integration with intelligent onboard decision-making, and extending current methods to more diverse environmental and acquisition conditions.

Satellites must increasingly be able to autonomously filter cloud-contaminated data before downlink, not only to conserve bandwidth and storage, but also to enable adaptive mission behavior—such as skipping low-quality scenes, re-tasking sensors, or prioritizing acquisitions based on content relevance. This requires not just lightweight models, but also robust AI frameworks optimized for real-world deployment constraints, including limited memory, power, and processing resources.

The HYPERBOLA project addresses this need by developing a scalable, AI-powered framework for real-time onboard cloud and shadow masking in hyperspectral imagery. Our approach combines the strengths of ViTs, GNNs, and CNN-Transformer hybrids to capture both spectral and spatial dependencies, while also incorporating inference optimizations such as structured pruning, 8-bit quantization, and TensorRT acceleration. Ultimately, the goal is to deliver high segmentation performance in a compact, energy-efficient architecture that can be deployed on embedded GPU systems such as the NVIDIA Jetson Xavier—laying the groundwork for more intelligent, autonomous, and scalable Earth observation missions.

In this paper, we present preliminary results focusing only on one candidate model: a hybrid CNN-Transformer architecture (Hybrid Model 2). This model demonstrates promising performance in segmenting dominant surface classes (Sea and Land), while also highlighting current limitations in detecting clouds. Future work will extend this evaluation to the full suite of HYPERBOLA models (ViTs, GNNs, and additional hybrid variants), incorporate optimizations such as structured pruning and 8-bit quantization, and benchmark real-time performance on embedded GPUs.

## 2. METHODS

### 2.1 Dataset

For training and evaluation, we used the publicly available HYPSONO-1 Sea-Land-Cloud dataset [5], a benchmark dataset tailored for hyperspectral cloud segmentation. The dataset includes 200 scenes acquired by the HYPSONO-1 satellite, each consisting of 108 contiguous spectral bands spanning from 400 to 2500 nm. Out of these, 38 scenes include pixel-wise annotations for three semantic classes: Sea, Land, and Cloud. After removing scenes with incomplete metadata or corrupted labels, we retained 27 high-quality labeled images. These were split into training (70%), validation (15%), and test (15%) subsets, ensuring a diverse distribution of geographical locations and atmospheric conditions.

Each hyperspectral cube was preprocessed by converting raw digital numbers into reflectance, followed by per-band normalization (zero-mean, unit variance) to stabilize training. Additionally, to enhance generalization under limited data, we applied data augmentation strategies, including horizontal and vertical flipping, random Gaussian noise injection in spectral bands, and gamma correction for contrast modulation. These techniques were designed to simulate variations in cloud texture, lighting, and atmospheric haze, as often encountered in real orbital conditions.

## 2.2 Model Architecture

For this preliminary study, we focused on a single architecture of interest: a Hybrid CNN-Transformer model (Hybrid Model 2). This design draws inspiration from U-Net-like encoder–decoder structures while incorporating transformer layers to capture global interactions across the scene. The encoder extracts multiscale spatial features using 2D convolutions, while the bottleneck processes flattened spatial tokens through transformer layers to model long-range dependencies. These global features are then fused during decoding to generate a full-resolution segmentation mask. By combining convolutional operations for local textural detail with transformer-based global context, the hybrid approach aims to balance efficiency with accuracy, particularly in challenging cases such as fragmented clouds or diffuse haze.

The model was trained using a weighted combination of cross-entropy loss and Dice loss to account for class imbalance between Sea, Land, and Cloud. Optimization was performed with the Adam optimizer, using early stopping and checkpointing based on validation Intersection-over-Union (IoU).

This architecture was selected as the focus of our preliminary results due to its promising trade-off between model complexity and segmentation performance. Future work will extend the evaluation to additional architectures, including Vision Transformers (ViTs) and Graph Neural Networks (GNNs), to further assess their potential in onboard hyperspectral cloud masking.

## 2.3 Onboard Optimization

To enable real-time inference on embedded satellite hardware, we are currently exploring a suite of model optimization strategies aimed at reducing memory usage, computational load, and power consumption—key constraints in onboard processing scenarios. These techniques are being evaluated for integration into the HYPERBOLA framework as part of future deployment phases.

One direction under investigation is structured pruning, where low-importance convolutional filters and transformer heads are iteratively removed to reduce model complexity. Structured pruning ensures tensor compatibility and efficient computation and has shown promise in maintaining accuracy with up to 30% reduction in parameter count in prior deep learning studies [9].

We are also testing post-training 8-bit quantization based on symmetric scaling, following the methodology of Jacob et al. [9]. This approach reduces the precision of weights and activations from 32-bit to 8-bit, potentially offering a 4× reduction in memory footprint while maintaining competitive inference performance. This is particularly relevant for models destined for edge AI devices with limited storage and memory bandwidth. These strategies are particularly relevant for the long-term evolution of the HYPERBOLA system toward scalable and modular satellite AI deployments.

# 3. RESULTS

## 3.1 Training performance – Preliminary Results

As part of the HYPERBOLA framework, we have defined four candidate architectures that capture different aspects of hyperspectral semantics (Table 1). These include a U-Net baseline, hybrid CNN-Transformer variants, and a Graph Neural Network extension.

Table 1. Comparison of HYPERBOLA models

Model Name	Base Architecture	GNN Module	Target Use
Hybrid Model 1	U-Net + ViT	No	Balanced
Hybrid Model 2	CNN + ViT	No	Efficiency
Model 3	U-Net (CNN)	No	Baseline
Hybrid Model 4	U-Net + ViT + GNN	Yes	Accuracy

At this stage, we report preliminary results focusing on Hybrid Model 2, which showed the most promising trade-off between model complexity and performance. On two representative test images, it achieved Dice scores of 0.7735 and 0.9329 for the Sea class, 0.7922 and 0.9346 for Land, and 0.3710 and 0.0 for Cloud, respectively. The corresponding IoU scores were 0.6307 and 0.8742 (Sea), 0.6560 and 0.8772 (Land), and 0.2278 and ~0.0 (Cloud). Notably, performance on the Cloud class varied considerably, with clear detection in Image 1 but no predictions in Image 2, likely due to either low contrast or near-absent cloud cover in that scene.

The overall average scores across both scenes were a Dice score of 0.5172 and IoU of 0.4851. These results confirm the model’s ability to accurately segment dominant surface classes (Sea and Land), while also underlining the current limitations in identifying thin, partially transparent, or sparse clouds.

As these results are preliminary, we are actively pursuing several enhancements to improve performance—particularly for the challenging Cloud class. Upcoming steps include switching from patch-wise to per-pixel normalization, which has already shown improved consistency during validation, and integrating custom loss functions with class-weighted tensors to better handle imbalanced class distributions. Furthermore, we plan to benchmark the model’s performance against the baseline performance previously established using these data and conventional CNNs [1]. This comparison will allow us to quantitatively assess the advantages of our hybrid approach in the context of cloud and cloud-shadow segmentation. Architecturally, we plan to refine the hybrid models by replacing simple skip-connection concatenation with spatial attention filters, aiming to enhance feature fusion across scales. While we also experimented with k-fold cross-validation, it yielded lower performance than single training runs, likely due to limited labeled data and scene variability. These targeted refinements are expected to substantially improve generalization and segmentation accuracy in future iterations of the HYPERBOLA framework.

### 3.2 Inference Evaluation – Preliminary Results

We evaluated Hybrid Model 2 (CNN + ViTs) on two representative test scenes from the HYPSO-1 dataset: Lake Volta (Ghana) and Bangladesh. These cases were selected to assess generalization across different geographical and spectral conditions.

In the Bangladesh scene, the model achieved strong segmentation for Sea and Land classes, with Dice scores of 0.9329 and 0.9346, and IoU scores of 0.8742 and 0.8772, respectively. However, it failed to identify any Cloud pixels in this scene, likely due to the absence or extremely subtle presence of cloud cover (Figure 1.).

In contrast, the Lake Volta scene (Figure 2.) demonstrated a more balanced class distribution, where the model successfully detected clouds with a Cloud Dice of 0.3710 and IoU of 0.2278, while also maintaining solid performance on Sea (Dice 0.7735, IoU 0.6307) and Land (Dice 0.7922, IoU 0.6560). The average across both scenes resulted in a Dice score of 0.5172 and IoU of 0.4851, reflecting a strong baseline for surface class segmentation and early-stage effectiveness in detecting cloud structures under certain conditions.




Figure 1. Ground truth (left) and predicted (right) image for Bangladesh.




Figure 2. Ground truth (left) and predicted (right) image for Lake Volta.

While these inference results align with the training behavior of Hybrid Model 2—especially its robustness on dominant surface classes—they also emphasize the need for refinements in detecting more complex and variable cloud formations. The enhancements outlined in Section 3.1 are expected to directly address these limitations and improve the model’s reliability in real-world onboard scenarios.

To provide a detailed breakdown of performance per class, we report precision, recall, Dice score, and Intersection over Union (IoU) for two representative test scenes: Lake Volta and Bangladesh. These results are presented in Table 2, highlighting both the model’s strengths in Sea and Land segmentation and its current limitations in Cloud detection.

Table 2. Metrics table (per class and overall) of the Hybrid Model 2 (CNN + ViTs).

Image	Metric	Precision	Recall	Dice Score	IoU Score
Lake Volta	Sea	0.6591	0.9359	0.7735	0.6307
Lake Volta	Land	0.6672	0.9750	0.7922	0.6560
Lake Volta	Cloud	0.9248	0.2321	0.3710	0.2278
Lake Volta	<b>Overall</b>	-	-	0.5993	0.4582
Bangladesh	Sea	0.9996	0.8745	0.9329	0.8742
Bangladesh	Land	0.8775	0.9996	0.9346	0.8772
Bangladesh	Cloud	0.0	0.0	0.0	0.0
Bangladesh	<b>Overall</b>	-	-	0.5172	0.4851

### 3.3 Limitations

Throughout the development and initial evaluation of the HYPERBOLA framework, several limitations were identified that currently constrain performance and generalizability.

First, the limited size and quality of the dataset posed a significant challenge. Although the HYPSON-1 dataset includes 38 labeled scenes, only 27 were usable due to missing or incomplete annotations. This reduction in effective training data restricts the model’s capacity to generalize across diverse atmospheric and surface conditions. In addition, the presence of

noisy or inconsistently labeled samples introduces uncertainty during training, potentially leading to overfitting or misclassification in ambiguous regions.

Another major constraint was class imbalance. Cloud pixels—particularly thin clouds or shadows—are often underrepresented relative to Sea and Land, which biases the learning process and reduces segmentation performance for the most critical class. This is compounded by the scarcity of publicly available satellite hyperspectral datasets with detailed cloud and shadow annotations, limiting opportunities for transfer learning or cross-domain validation.

From a modeling perspective, we observed that higher-performing architectures are often more computationally intensive, raising concerns about their feasibility for onboard deployment. While Hybrid Model 2 currently offers a promising trade-off between accuracy and efficiency, more complex configurations like Hybrid Model 4—featuring Graph Neural Networks—may not be suitable for real-time, energy-constrained platforms.

Lastly, the relationship between model complexity and operational value remains an open question. In some cases, simpler architectures may outperform advanced models when robustness, interpretability, or latency are prioritized. Future evaluations will need to balance these trade-offs carefully, especially as the project progresses toward hardware-aware deployment scenarios.

#### 4. DISCUSSION

The preliminary results from the HYPERBOLA framework highlight both the potential and the challenges of deploying advanced AI models for onboard cloud and cloud-shadow segmentation in hyperspectral imagery. Our experiments with Hybrid Model 2 (CNN + ViT) demonstrated promising segmentation performance across dominant surface classes (Sea and Land), with good generalization between geographically distinct scenes. This supports the hypothesis that transformer-based architectures can effectively capture long-range spectral dependencies, even under constrained training data conditions. However, the variability in cloud detection performance highlights the sensitivity of the model to class imbalance, annotation noise, and subtle spectral signatures—issues that are intrinsic to many remote sensing datasets. The Cloud class, in particular, proved to be the most challenging, as its detection often depends on thin, partially transparent, or sparse clouds that are difficult to identify with conventional approaches. This suggests that improvements in data preprocessing, loss design, and normalization strategies may yield larger gains than merely increasing model complexity. Our proposed shift to per-pixel normalization and class-weighted loss functions reflects this understanding and aligns with prior findings in imbalanced learning. At this stage, our focus is on establishing a strong baseline with Hybrid Model 2. Future work will extend the evaluation to additional architectures in the HYPERBOLA framework (ViTs, GNNs, and hybrid variants) to systematically explore trade-offs between accuracy, robustness, and computational efficiency for onboard deployment. It is also evident that performance evaluation must go beyond global metrics such as Dice scores and IoU. Scene-specific conditions—such as illumination, cloud density, and surface texture—play a significant role in model behavior, as demonstrated by the performance discrepancy between the Lake Volta and Bangladesh scenes. The model performed better in the Lake Volta scene, where cloud coverage was more prominent, while its performance in the Bangladesh scene suffered due to the near-absence of cloud cover or extremely subtle cloud types. Future benchmarks should incorporate a more diverse set of annotated hyperspectral scenes and assess robustness under varying orbital and atmospheric conditions. The broader implication of this work lies in its contribution to enabling more intelligent and autonomous satellite missions. Accurate onboard cloud masking is essential for reducing bandwidth requirements, prioritizing downlink content, and supporting real-time decision-making. By automatically filtering out low-quality, cloud-contaminated data before downlink, satellites can save valuable resources and transmit only the most relevant information to ground stations. This can also support mission flexibility, such as re-tasking sensors or triggering secondary analytics pipelines in response to specific detected features. By demonstrating that transformer-based models can operate at high fidelity within realistic onboard constraints, HYPERBOLA begins to lay the groundwork for a new class of hyperspectral-aware space systems that are more autonomous, scalable, and resource-efficient. Moving forward, the focus will shift toward integrating the optimization strategies outlined in Section 2.3, including structured pruning, 8-bit quantization, and deployment through TensorRT acceleration. Additionally, expanding the dataset through semi-supervised labeling or domain adaptation techniques will help improve the model's performance in regions with limited annotated data. Further testing on embedded platforms such as the NVIDIA Jetson AGX Xavier will help bridge the current gap between experimental prototypes and flight-ready onboard AI systems. These steps will ensure that HYPERBOLA can operate effectively in real-world spaceborne scenarios and serve as a scalable, reliable solution for hyperspectral data analysis in Earth Observation missions.

## 5. CONCLUSIONS

In this work, we presented **preliminary results** from the HYPERBOLA project, an AI framework for real-time cloud masking in hyperspectral imagery onboard satellite platforms. At this stage, we focused on a single candidate architecture, **Hybrid Model 2 (CNN + ViT)**, which demonstrated strong segmentation performance for Sea and Land classes and highlighted current challenges in detecting thin or sparse clouds. These findings establish an early baseline for transformer-based approaches to onboard hyperspectral cloud segmentation. The variability in cloud detection performance underscores the importance of continued refinement, particularly in handling class imbalance and improving robustness under complex atmospheric conditions. Planned improvements include per-pixel normalization, class-weighted loss functions, structured pruning, and quantization, all aimed at enhancing both segmentation accuracy and onboard deployability. Future work will expand the evaluation to the broader set of HYPERBOLA architectures—including Vision Transformers, Graph Neural Networks, and additional hybrid variants—in order to systematically explore trade-offs between accuracy, robustness, and computational efficiency. These studies, combined with hardware-aware optimization and deployment on embedded platforms such as the NVIDIA Jetson Xavier, will be essential steps toward flight-ready implementations. Ultimately, HYPERBOLA aims to enable a new generation of intelligent, autonomous Earth Observation systems. By analyzing and filtering hyperspectral data directly in orbit, satellites can reduce transmission loads, prioritize high-value observations, and support more adaptive and efficient mission operations.

## 6. ACKNOWLEDGMENTS

This work is supported by the “ENFIELD: European Lighthouse to Manifest Trustworthy and Green AI” project funded by the European Union’s HORIZON Research and Innovation Programme under Grant Agreement No. 101120657. We thank the NTNU SmallSat Lab for the HYPSON-1 dataset. We are also grateful to our mentors from the Big Earth Data Analytics Department at the Eratosthenes Centre of Excellence — in particular, Dr. Michalis Mavrouniotis and Stelios Neophytides— for their valuable feedback, scientific guidance, and continued support throughout this research.

## 7. REFERENCES

- [1] J. A. Justo, A. Ghiță, D. Kováč, J. L. Garrett, M.-I. Georgescu, J. Gonzalez-Llorente, R. T. Ionescu, and T. A. Johansen, “Semantic Segmentation in Satellite Hyperspectral Imagery by Deep Learning,” *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 18, 273–293 (27 Feb 2025). <https://doi.org/10.1109/JSTARS.2024.3487360>
- [2] G. Camps-Valls, D. Tuia, L. Bruzzone, and J. A. Benediktsson, “Advances in Hyperspectral Image Classification: Earth Monitoring with Statistical Learning Methods,” *IEEE Signal Process. Mag.* 31(1), 45–54 (1 Jan 2014). <https://doi.org/10.1109/MSP.2013.2279179>
- [3] Kumar, V., Singh, R. S., Rambabu, M., and Dua, Y., “Deep learning for hyperspectral image classification: A survey,” *Comput. Sci. Rev.* 53(8), 100658 (1 Aug 2024). <https://doi.org/10.1016/j.cosrev.2024.100658>
- [4] Li, S., Song, W., Fang, L., Chen, Y., Ghamisi, P., and Benediktsson, J. A., “Deep learning for hyperspectral image classification: An overview,” *IEEE Trans. Geosci. Remote Sens.* 57(9), 6690–6709 (20 Mar 2019). <https://doi.org/10.1109/TGRS.2019.2907932>
- [5] Ghasemi, N., Justo, J. A., Celesti, M., Despoisse, L., and Nieke, J., “Onboard processing of hyperspectral imagery: Deep learning advancements, methodologies, challenges and emerging trends,” *arXiv preprint arXiv:2404.06526v1* (9 Apr 2024); <https://arxiv.org/abs/2404.06526v1>
- [6] Justo, J. A., Garrett, J., Langer, D. D., Henriksen, M. B., Ionescu, R. T., and Johansen, T. A., “An Open Hyperspectral Dataset with Sea-Land-Cloud Ground-Truth from the HYPSON-1 Satellite,” *arXiv preprint arXiv:2308.13679v2* (3 Sep 2023); <https://arxiv.org/abs/2308.13679>
- [7] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, et al., “An Image is Worth 16×16 Words: Transformers for Image Recognition at Scale,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 26 Apr 2021; <https://arxiv.org/abs/2010.11929>
- [8] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. Z. Yu, “A Comprehensive Survey on Graph Neural Networks,” *IEEE Trans. Neural Netw. Learn. Syst.* 32(1), 4–24 (15 Jan 2021). <https://doi.org/10.1109/TNNLS.2020.2978386>
- [9] B. Jacob, S. Kligys, B. Chen, et al., “Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Salt Lake City, UT, USA, 18–22 Jun 2018. <http://dx.doi.org/10.1109/CVPR.2018.00286>