

Leveraging Deep Learning for High-Resolution Optical Satellite Imagery From Low-Cost Small Satellite Platforms

Valentino Constantinou¹, Mark Hoffmann¹, Matthew Paterson¹, Ali Mezher¹, Brian Pak¹, Alexander Pertica¹, and Emily Milne

Abstract—The expansion of small satellite networks in earth’s orbit has resulted in a plethora of earth optical imagery available to the civil, defense, and commercial sectors. Small satellites (less than 1000 kg in mass) and their constellations can be delivered rapidly and at low cost and are more difficult to target by adversaries—a key consideration in the defense industry. Yet, small satellite size constraints often result in reduced payload capacity, reduced power capacity, or loss of redundancy. Traditionally, the cost of an optical telescope on board a satellite scales at roughly the square of the aperture, meaning that it costs four times as much to double the resolution of the imaging hardware. However, deep learning has shown considerable success in the areas of super-resolution and enhancing the pixel resolution of optical imagery. These deep learning methods have the potential to provide optical resolution capabilities rivaling larger satellites and their telescopes, while maintaining the benefits of small satellites—smaller physical size (which lowers launch vehicle costs and provides a basis for large constellations), reduced manufacturing time, and lower manufacturing costs. By providing low-cost small satellite platforms with the same capabilities as larger satellites, the cost for high-resolution in-orbit optical imagery is reduced alongside time to orbit. In this work, we detail a deep-learning-based approach, which improves optical satellite imagery to five times the original pixel-based resolution without the need or expense of increasing the capabilities of the imager through larger telescope apertures. The approach—demonstrated on Terran Orbital’s GEOSTare SV2 mission imagery—is generally applicable to any optical satellite image and is agnostic to the mission, satellite manufacturer, optical payload specifications, or data source. This capability provides a basis for small satellite missions and constellations—and their optical payloads—to rival the native hardware-based resolutions available through larger satellites with wider telescope apertures at a significantly reduced cost.

Index Terms—Deep learning, deep neural networks (DNNs), optical imagery, remote sensing, satellite, super-resolution (SR).

I. INTRODUCTION

THE advent of commercially available high-volume microelectronics has enabled the growth of small satellites in the aerospace industry, particularly in low earth orbit (LEO). Small satellites are generally considered those to be less than

1000 kg in mass [1]. The hallmark of today’s modern small satellites is the adoption of consumer technologies—combined with rapid development and production—by small agile teams more closely aligned with information technology industry management models [2]. Typically, small satellite design and manufacturing teams exhibit highly innovative and motivated technical staff, good communication with teams in close proximity, well-defined mission objectives, short time scales for delivery and launch of operations, failure-resistant system architectures, and an objective to design to cost [1].

Robert “Bob” Twiggs and Dr. Jordi Puig-Suari introduced the small satellite CubeSat standard in 1999 [1], with Puig-Suari and Scott MacGillivray later cofounding Tyvak Nanosystems—now Terran Orbital—in 2011. Since its introduction, the CubeSat has been the gold standard in small satellite design for both research and commercial platforms. In 2013, the first commercial CubeSats with propulsion and precision control were introduced, followed by the first CubeSat constellation in 2015, the first CubeSats for controlled orbit changes and rendezvous in orbit in 2017, the first CubeSats leaving Earth orbit in 2018 (the MARCO CubeSats supporting the National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory’s InSight mission), the first CubeSat launch with a radar for NASA Jet Propulsion Laboratory in 2018 (RainCube), and the first in-orbit nanosatellite docking mission in 2022 (CPOD). Since their introduction, CubeSats and their form factors have become the predominant architecture for designing, developing, engineering, and launching small satellites in earth orbit and beyond.

Recently, small satellite constellations have been gaining relevance in the defense and commercial sectors. The Space Development Agency (SDA) is a United States Space Force (USSF) unit tasked with deploying disruptive space technology to LEO, geostationary equatorial orbit (GEO), and other space environments. Military officials from the USSF and related agencies have publicly touted the benefits of a globally distributed network of small satellites, with both the difficulty of targeting and destroying a small satellite versus more traditional large satellite platforms. On April 5, 2023, SDA director Derek Tournear stated “We’ll have hundreds and hundreds of these satellites up there [with the Proliferated Warfighter Space Architecture]. It will cost more to shoot down a single satellite than it will cost to build that single satellite. We just completely

Manuscript received 13 November 2023; revised 9 January 2024; accepted 3 February 2024. Date of publication 14 February 2024; date of current version 15 March 2024. (Corresponding author: Valentino Constantinou.)

The authors are with the Terran Orbital Corporation, Irvine, CA 92618 USA (e-mail: valentino.constantinou@terranorbital.com).

Digital Object Identifier 10.1109/JSTARS.2024.3365417




Fig. 1. GEOSTare SV2 satellite bus and payload. GEOSTare SV2 is a 6U cubesat about 36 cm tall and 120 cm long, containing two optical telescopes—V3 (wide field of view) and V4 (narrow field of view).

changed that value equation.” Within the United States defense sector, there is an ever greater emphasis on constellations of small satellites, which provide increased redundancy, reliability, and operational stability compared to more traditional larger satellite platforms. The United States Department of Defense’s Operationally Responsive Space concept called for the rapid development and launch of spacecraft to augment or replace existing spacecraft, with the goals to develop new small satellites using standardized buses and plug-and-play architectures to dramatically reduce development times and costs [1]. In the commercial sector, companies like SpaceX (with their Starlink constellation) and Planet (with their constellation of “Doves”) are proliferating thousands of satellites into earth orbit with extensive spatiotemporal coverage. As of October 1, 2023, there have been over 4200 SpaceX Starlink satellites, 150 Planet Dove satellites, and 630 OneWeb satellites in LEO. Rivada Networks plans a constellation of over 600 satellites for encrypted optical communications in LEO [3].

Regardless of the payload or capability facilitated through these constellations in earth orbit, the trend is clear—constellations of small satellites provide reduced manufacturing costs, reduced launch costs, greater redundancy, greater resiliency against attack by adversaries, and a vastly improved ability to cover all areas of the globe from earth orbit—the real utility emerges through the formation of low-cost small satellite constellations [1]. These satellite constellations also provide the ability to upgrade the constellations over time with new bus or payload capabilities, as demonstrated with SpaceX’s Starlink V1, V1.5, and V2 satellites or the SDA’s Proliferated Warfighter Space Architecture—each new version or tranche replacing or augmenting existing satellites with upgraded platforms and capabilities (Fig. 1).

Terran Orbital’s GEOSTare SV2 satellite—developed together with the Lawrence Livermore National Laboratory (LLNL)—launched on May 15, 2021 on board a Space-X Falcon-9 launch vehicle [4]. The nearly 36-cm-tall 120-cm-long GEOSTare SV2 is the latest development of a collaboration between LLNL and Terran Orbital to demonstrate the feasibility of using these small systems for space domain awareness (monitoring and tracking space debris), astronomical observation, and ground imaging. In the first six months of operation, GEOSTare SV2 captured over 29 000 images of earth and space at a maximum 2-m ground sampling distance (GSD). GEOSTare SV2’s carries a

NVIDIA TX2i payload computer. The inclusion of an NVIDIA graphical processing unit (GPU) onboard GEOSTare SV2 provides an opportunity for edge computing of artificial intelligence (AI) solutions from earth orbit, although there are cases where machine and deep learning capabilities are best deployed “on the ground” to limit the impact of in-orbit operations.

Small satellites, such as GEOSTare SV2, are attractive to many civil, defense, and commercial customers due to their reduced cost, reduced assembly time, and easier integration with launch vehicles due to their sizing and available fairing capacity. However, small satellite platforms, such as GEOSTare SV2, are subject to various limitations, which reduce the capabilities of onboard payloads. Tradeoffs are made between the size of the satellite and the available resources that can be provided on board the platform. In the case of GEOSTare SV2, while its optical imaging payloads represent an industry advancement, physical sizing limitations result in reduced telescope apertures and thus GSD. A major challenge of high-resolution imaging missions on small spacecraft is the thermal stability on the imaging system and satellite bus needed to provide sharp imagery [5].

AI approaches, such as deep neural networks (DNNs), present an opportunity to “do more with less,” bringing capabilities traditionally associated with large satellites and large-diameter telescopes to small satellite platforms. By integrating DNN capabilities, small satellite operators can offer data products rivaling those offered by larger satellites and payloads at a fraction of the hardware cost and time to orbit associated with larger telescope diameters or satellites. Super-resolution (SR) techniques find many applications across industries, including satellite imagery [6]. SR is effectively a one-to-many inference problem from a low-resolution space to a high-resolution space, of which determining the correct solution is not trivial [7]. Nevertheless, the flexibility of DNNs and their capability to learn highly nonlinear patterns has shown to attain excellent image reconstruction accuracy compared to handcrafted models [8], [9], [10]. More recently, single-image super resolution (SISR) techniques have sought to learn implicit redundancy that is present in data to recover missing high-resolution information from a single low-resolution sample [7], [11]. A comprehensive review of SISR techniques can be found in [12]. High-resolution imagery exhibits greater texture detail necessary to separate buildings from their backgrounds, a critical consideration in producing high-quality building footprints [13], essential for urban planning, disaster response, and population estimation from earth orbit [14], [15], [16]. SR helps to enhance the distinction between backgrounds and buildings in the resulting images [17].

While the cost of an optical telescope on board a satellite scales at roughly the square of the aperture, meaning it that costs four times as much to double the resolution using hardware alone (millions or tens of millions of USD), doing the equivalent through SR using DNNs is achievable in hundreds of thousands of dollars and at reduced time scales (months versus years of development). Furthermore, these techniques can be deployed using existing ground station infrastructure or by using edge computing in-flight (on board the spacecraft) such as with GEOSTare SV2’s NVIDIA TX2i or on the ground using on-premises or cloud computing infrastructure (on-demand),



Fig. 2. Side-by-side comparison of iSAID-derived training data, with the low-resolution input image (242×242 pixels) on the left and the high-resolution target image (1210×1210 pixels) on the right.

offering greater flexibility in satellite payload offerings. Computationally intensive model training can be accomplished on the ground, with fully trained and optimized models then uploaded and deployed to the spacecraft for onboard processing of mission data [18].

II. METHODS

In the subsequent sections, we detail our SISR image improvement methodology for optical satellite imagery of the earth surface.

In summary, during both model training and inference, input images are tiled into 64 square images. For model training, the 64 square images are kept as the high-resolution targets and are resized to 1210×1210 pixels. Low-resolution versions of the same 64 square images are used as the low-resolution inputs, resized to 242×242 pixels (Fig. 2). Together, these input and target images comprise training data. This methodology leverages a publicly available dataset for training the well-known ResNet (34-layer) [19] architecture as a backbone for a U-Net [20]. After inference, tiled sections of each image are stitched together to result in a reconstructed high-resolution image.

We validate the performance of the model both quantitatively and qualitatively. Using GEOStare SV2 optical satellite imagery for inference, we upscale the pixel resolution of the satellite imagery by five times. Errors between the inferred high-resolution image and the original image are illustrated (for both full images and tiles of images). Side-by-side comparisons are made between the native-resolution GEOStare SV2 imagery and the upscaled high-resolution versions in addition to the quantitative evaluation.

A. Training Data

Large-scale datasets, such as ImageNet [21], PASCAL-VOC [22], and MSCOCO [23], have enabled the training of data-hungry DNNs, bringing unprecedented performance in scene understanding tasks such as image classification [19], [24], object detection [25], [26], and instance segmentation [27], [28]. However, these datasets contain terrestrial natural scenes where objects are appeared in an upward orientation and do not transfer well to aerial imagery such as those from in-orbit imaging satellites or unmanned aerial vehicles [29]. In aerial imagery, objects occur in high density, large aspect ratios, and with variation in scale, shape, and orientation [30].

A training dataset is created using publicly available large-scale dataset for instance segmentation in aerial images

(iSAID) [30]. The images in the iSAID data are collected from Google Earth, with most of the images captured by satellites JL-1 (0.75-m GSD) and GF-2 (0.81-m GSD). The distinctive characteristics of iSAID are the following: 1) large number of images with high spatial resolution; 2) 15 important and commonly occurring categories; 3) large number of instances per category; 4) large count of labeled instances per image, which might help in learning contextual information; 5) huge object scale variation, containing small, medium, and large objects, often within the same image; 6) imbalanced and uneven distribution of objects with varying orientation within images, depicting real-life aerial conditions; 7) several small size objects, with ambiguous appearance, can only be resolved with contextual reasoning; and 8) precise instance-level annotations carried out by professional annotators, cross-checked and validated by expert annotators complying with well-defined guidelines.

The iSAID data contain 937 images used for the validation of image segmentation models, which we repurposed for training our SR optical satellite ground imagery model. Each source image from the iSAID dataset is tiled into 64 square images in both low-resolution input images (242×242 pixels) and high-resolution target images (1210×1210), which are five times the resolution of the input images, together creating low-resolution and high-resolution pairs. Both sets thus contain 59 968 images (937 iSAID images tiled into 64 pieces). Together, these two sets of images comprise the training data. By training on the model on three-channel red–green–blue (RGB) images, the resulting trained model can be applied to a wide variety of input data for inference, including 2-D representations of synthetic aperture radar, multispectral and hyperspectral imagery. By training the model on imagery with a smaller GSD compared to GEOStare SV2 [75 cm (JL-1) and 80 cm (GF-2) compared to 2 m (GEOStare SV2)], we provide and encode the model with information regarding patterns and textures in optical satellite imagery that is unavailable or hard to distinguish in the native GEOStare SV2 imagery (and its resolution).

B. GEOStare SV2 V3 and V4 Optical Telescopes

The GEOStare SV2 spacecraft contains the GEOStare2 optical imaging payload. The payload holds two compact Livermore-developed monolithic telescopes (MonoTeles), along with associated sensors, cameras, and electronics. One MonoTele has a high-resolution narrow field of view for earth observation (V4); the other has a high-sensitivity wide field of view (V3) for space domain awareness and astronomy applications. The MonoTeles are developed from a single monolithic fused silica slab with reflective coatings on each end. This design reduces the effects of vibration at launch and eliminates the need for adjusting the sensors on-orbit.

C. GEOStare SV2 Optical Imagery

The images produced from the V3 and V4 telescopes are 1936×1216 and 3096×2080 pixels in resolution, respectively, with a GSD of approximately 2 m for the V4 telescope imagery and 10.4 m for the V3 telescope. All the GEOStare SV2 images are black-and-white single-channel images and exhibit various illumination conditions and viewing angles in the satellite images due to the 53° inclined orbit, which is not



Fig. 3. Image from GEOStare SV2's V3 (wide field of view) telescope.

in sun-synchronous orbit (SSO). This orbit is not common in commercial satellite imagery, which is predominantly produced by satellites and telescopes in SSO. The satellite can be rapidly retasked for different image locations and times to capture these artifacts, supplemented by an extensive ground station network. When considered under the National Imagery Interpretability Rating Scale (NIIRS) [31], GEOStare SV2 imagery falls within rating level 3—“It is possible to detect large area contour ploughing; individual houses in residential areas; trains or strings of rolling stock (but not individual wagons); identify inland waterways navigable by barges; distinguish between natural forest and orchards.”

The GEOStare SV2 optical imagery from both the V3 and V4 telescopes are used to both quantitatively and qualitatively assess the performance of the trained model. To assess the imagery quantitatively, the native-resolution 242×242 tiled images and the high-resolution 1210×1210 are directly compared by resizing the latter to the native 242×242 pixel resolution. More specifically, each 1210×1210 high-resolution image tile is resized to the 242×242 pixel resolution inherent in the native-resolution imagery. Each of these tiles is then compared with the corresponding native-resolution tile, with the absolute differences between each pixel calculated, illustrated, and captured. In Fig. 8, the native-resolution (left), high-resolution (center), and differences between the two (right) are illustrated for one example set of tiles. The mean error per pixel is also captured by averaging the absolute error across all of the native-resolution and corresponding high-resolution image tiles, shown and discussed in Fig. 10.

To assess the imagery qualitatively, we directly compare cropped areas (i.e., “zoomed-in” areas) of both native- and high-resolution GEOStare SV2 imagery against the NIIRS rating level system across a wide set of examples, including urban city centers, airports, ports, agricultural fields, natural landscapes, and other geographic settings. Two examples from urban city centers are illustrated and discussed in Section III.

D. Deep Learning Model Architecture

We employ an SISR method to reconstruct high-resolution optical ground imagery from satellites from single low-resolution inputs, without the need for a sequence of images. More specifically, we utilize the well-known ResNet (34 layer) [19] architecture as a backbone for a U-Net. A U-Net is a convolutional



Fig. 4. Image from GEOStare SV2's V4 (narrow field of view) telescope.

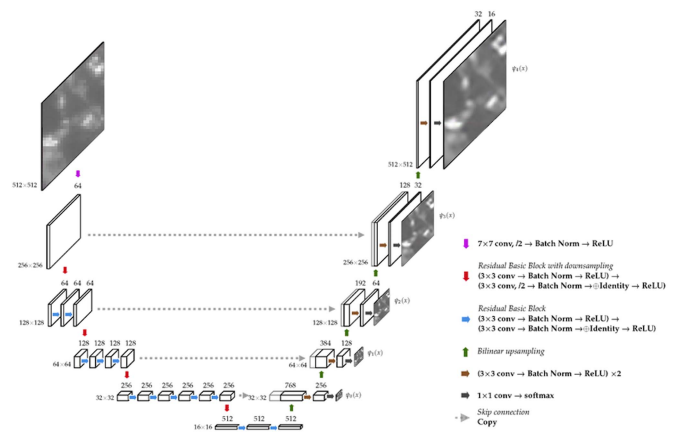


Fig. 5. Illustration of the U-Net architecture with a ResNet (34-layer) backbone, using a native-resolution (left) and high-resolution (right) image of Lower Manhattan as an example.




Fig. 6. Training and validation loss curves of the SR optical imagery model over 11 epochs of model training.

neural network (CNN) that was originally developed for image segmentation at the Department of Computer Science, University of Freiburg. The architecture of U-Net models consists of a contracting path and an expansive path, which gives it the U-shaped architecture (Fig. 5). The contracting path is a typical convolutional network that consists of repeated application of convolutions, each followed by rectified linear unit activation functions and a max-pooling operation. During the contraction, the spatial information is reduced, while feature information




Fig. 7. Side-by-side comparison of iSAID-derived training data, with the low-resolution input image (242×242 pixels) on the left, the high-resolution target image (1210×1210 pixels) in the center, and the high-resolution inferred (predicted) image (1210×1210 pixels) on the right.

is increased. The expansive pathway combines the feature and spatial information through a sequence of upconvolutions and concatenations with high-resolution features from the contracting path [20].

U-net architectures exhibit higher computational complexity and cost compared with other model architectures commonly used in SISR or segmentation applications. Due to the skip connections present in the U-net architecture and the large number of parameters, U-nets are more prone to overfitting and require additional computational cost or time to train. However, the use of skip connections allows U-nets to incorporate both high- and low-level features from an input image, typically improving the output from the trained models.

We train the U-Net model with the backbone of a residual neural network, or ResNet, utilizing a ResNet (34-layer) encoder and decoder (parameters noted in Table III). ResNet models are other CNNs made up of a series of residual blocks with skip connections, differentiating ResNets from other CNNs [19]. Since CNNs can be substantially deeper, more accurate, and more efficient to train if they contain shorter connections between layers close to the input and those close to the output, ResNet architectures add cross connections between the layers of the network allowing large sections of the network to be skipped if needed. When using only a U-Net, predictions tend to lack fine detail, but utilizing a ResNet backbone with cross or skip connections can address this inherent limitation, resulting in improved high-resolution imagery with the ability to capture fine details necessary to improve the quality of optical ground imagery.

E. Loss Function

We utilize the mean square error (MSE) to calculate the difference between the pixels predicted by the SISR model and the target images, defined as

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2$$

where Y_i and \hat{Y}_i are the target and predicted high-resolution images of input image i and N is the number of training data pairs. Other loss functions such as the mean absolute error or binary accuracy have been implemented in related literature, but were not explored in the course of this work.

F. Model Training Computing Resources

We utilized an Amazon Web Services (AWS) EC2 instance for model training, specifically the g4dn.12xlarge instance type with four NVIDIA Tesla T4 GPUs. Each NVIDIA Tesla T4 contains 2560 CUDA cores, 320 Turing Tensor Cores, and 16 GB of GDDR6 random access memory (RAM). We execute and accelerate model training by distributing the batches across each GPU, training four batches in parallel for each iteration in each epoch. The Python [32] programming language was used, Numpy [33], PyTorch [34], and FastAI [35] for model training and inference, and the Matplotlib [36], Seaborn [37], and Pillow libraries for visualization. Model training takes approximately 4.5 days using our training set derived from iSAID. Using our AWS inference infrastructure, improving the resolution of a GEOSTare SV2 image takes approximately 7 min using only the central processing unit.

III. RESULTS

A. Model Evaluation

The training and validation loss curves from model training are shown in Fig. 6. Early stopping and learning rate reduction on plateauing are implemented as part of the model training process. Close examination of the model training and validation loss curves shows a slight tendency toward overfitting after epoch 10. Future implementations of this capability could stand to more aggressively implement early stopping or a reduction in the learning rate as progress plateaus, further increasing modeling accuracy and performance, at the expense of increased model training times.

B. Quantitative Evaluation of High-Resolution Imagery

For quantitative evaluation of model performance, we evaluate the absolute error between each pixel across the native-resolution inputs and the high-resolution model outputs. More specifically, we evaluate the model on 250 randomly sampled ground images from GEOSTare SV2, resulting in 16000 tiled images (242×242 pixels each for native resolution and 1210×1210 pixels each for high-resolution from the model output). For each pair of native- and high-resolution imagery, we resize the high-resolution 1210×1210 pixel images back to the original 242×242 size and then calculate the absolute error of each pixel between the native- and high-resolution images. A result of this comparison is shown in Fig. 8. We then illustrate the errors in the fully reconstructed images in addition to the tiles, as shown in Fig. 9.

In addition, we calculate and illustrate the mean (average) error across all the pixels for each tile to determine the distribution of errors across pixels in order to examine where in each image tile the model exhibits the most error. Intuitively, the model exhibits the greatest error at the borders of each tiled image, where less information is present for image improvement due to a decrease in the amount of neighboring pixels (shown in Fig. 10).




Fig. 8. 242 × 242 native pixel resolution image tile (left), a 1210 × 1210 high-resolution image tile (center), and errors (right) between the two image tiles. Note that the greatest error manifests in this example surrounding complex features, such as cloud boundaries. There is little to no error in the area within the cloud itself.




Fig. 9. Reconstructed image with errors overlaid on top of the optical image. Note that the greatest error manifests in the borders of the tiles and in complex features, such as building boundaries or topographical changes. There is little to no error in the area within less complex features such as the water surface in the canals or accompanying lake.

C. Qualitative Comparison of Native- and High-Resolution Imagery

In addition to a quantitative evaluation of the output imagery, a qualitative assessment of the resulting reconstructed imagery is performed as in [17]. A goal of our methodology is to produce GEOSTare SV2 imagery at NIIRS rating level 4 from the existing rating level 3. In NIIRS rating level 4, “It is possible to identify farm buildings as barns, silos or residences; detect basketball or tennis courts in urban areas; identify individual tracks, rail pairs and control towers; detect jeep trails through grassland.” In order to assess whether NIIRS rating level 4 has been achieved, we examine the high-resolution model outputs against the native-resolution ground imagery from GEOSTare SV2 at microscales by cropping representative sections of each image.

We illustrate two examples from GEOSTare SV2 imagery from New York City in the United States and Kuala Lumpur in Malaysia. The high-resolution improved image of New York City with a bounding box over the cropped area is shown in Fig. 11, with the microscale cropped section illustrating




Fig. 10. Heatmap of the mean (average) absolute errors across all of the tiled images. With the exception of the tiled image boundaries, the error is (generally) normally distributed across each image tile.



Fig. 11. Improved high-resolution optical satellite image of New York City, including Lower Manhattan. The area indicated in the green box is expanded and shown in Fig. 12.

lower Manhattan shown in Fig. 12. The comparison between the native-resolution GEOSTare SV2 image (left) and high-resolution image (right) illustrates the improvements to the source image made with the technique outlined in this article, mainly: 1) additional clarity around building and wharf edges; 2) the improved ability to determine building characteristics such as shape and height; and 3) an improved ability to determine building footprints from adjacent roadways. Whereas the original imagery proves difficult to identify building boundaries and footprints and thus building types, the high-resolution imagery allows viewers to better identify individual buildings and their structures.

Similarly, we illustrate an example of a moving train in Kuala Lumpur, Malaysia. The high-resolution improved image of Kuala Lumpur with a bounding box over the cropped area




Fig. 12. Lower Manhattan in New York City, with the native-resolution GEOSTare SV2 image shown on the left and the high-resolution model result shown on the right. Note the dramatically improved ability to interpret and estimate building footprints in the improved image.

TABLE I
GENERAL CLASSIFICATION OF SATELLITES BY MASS [1]

Classes of Satellites by Mass (in kilograms, kg)	
Large satellite	≥ 1000 kg
Small satellite	500–999.99 kg
Mini satellite	100–500 kg
Micro satellite	10–100 kg
Nano satellite	1–10 kg
Pico satellite	0.11–1 kg
Femto satellite	≤ 0.1 kg

TABLE II
TECHNICAL SPECIFICATIONS OF THE V3 AND V4 TELESCOPES AND THE IMAGERS

Telescope and Imager Specifications	
Imager	VIS-Wide
Telescope	V3
Aperture (mm)	85
Focal length (mm)	306
F/number	3.6
Focus range	4km to infinity
X, Y pixels	1936x1216
GSD at 544km	10.4m
Coverage at 544km	255.7km ²
Imager	VIS-Narrow
Telescope	V4
Aperture (mm)	85
Focal length (mm)	650
F/number	8.0
Focus range	2km to infinity
X, Y pixels	3096x2080
GSD at 544km	2.0m
Coverage at 544km	26.0km ²



Fig. 13. Improved high-resolution optical satellite image of Kuala Lumpur, Malaysia. The area indicated in the green box is expanded and shown in Fig. 14.

is shown in Fig. 13, with the microscale cropped section illustrating the moving train shown in Fig. 14. The comparison between the native-resolution GEOSTare SV2 image (left) and

high-resolution image (right) illustrates the improvements to the source image made with the technique outlined in this article, mainly: 1) additional clarity around building footprints and transportation networks; 2) improved clarity around moving objects in the image such as the train; and 3) an ability to determine that the train contains two moving cars, shown with a small dark “break” between each rail car. Whereas the train is more difficult to identify in the original image, the high-resolution image clearly shows the moving train.



Fig. 14. Moving train captured in Kuala Lumpur, Malaysia, with the native-resolution GEOSTare SV2 image shown on the left and the high-resolution model result shown on the right. Note the dramatically improved ability; see building details and the moving train in the upscaled satellite image.

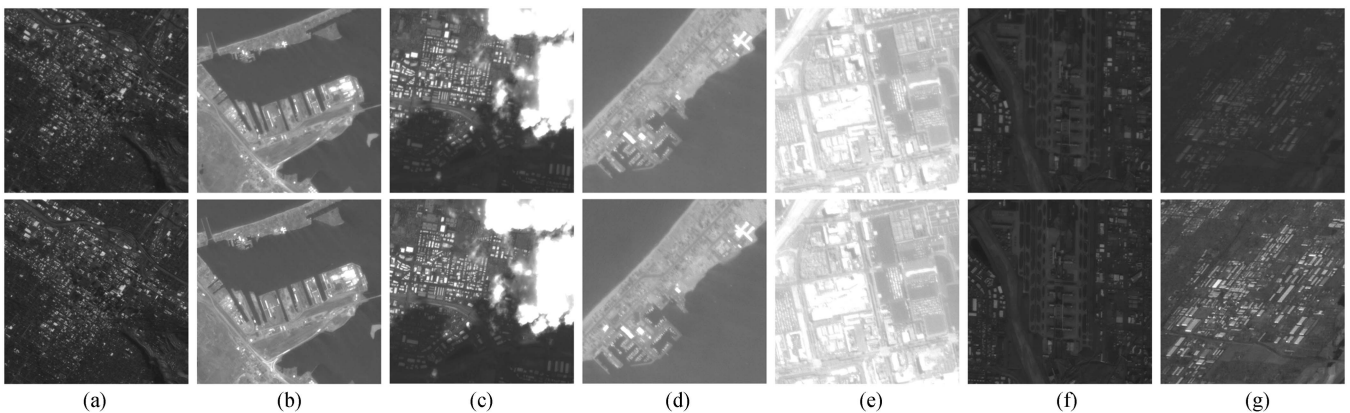


Fig. 15. Input images (top) and the results (bottom) of applying the SISR approach under various imaging conditions. (a) Dawn/dusk. (b) Mid-day. (c) Morning/afternoon with partial cloud cover. (d) With noise from Bayer filtering. (e) Overexposure. (f) Underexposure. (g) Oblique angle imagery.

TABLE III
PARAMETERS OF THE MODEL USED IN THE EXPERIMENTS

Model Parameters	
Architecture	U-Net with ResNet (34-layer) backbone
Batch size	1
Optimizer	Adam
Beginning learning rate	0.00011
Loss function	MSE
Input image size (in pixels)	242x242
Output image size (in pixels)	1210x1210

D. Robustness to Imaging Variabilities

It is well known that multi- and hyperspectral data can suffer from noise, degradation, and other variabilities introduced through imaging [38]. With respect to three-channel optical imaging, variabilities in the imaging process can be introduced due to environmental conditions on board the satellite (the

thermal properties of the imager and satellite at the time of capture), environmental conditions between the satellite and the ground (weather or the satellite's position relative to the sun), and errant noise introduced through errors in onboard or ground data processing. In Fig. 15, we highlight the effect of various imaging variabilities on the results of our SISR approach. We can see from the results in the figure that the SISR approach performs well under various lighting conditions and other variabilities such as noise from Bayer filtering present in the V4 telescope.

E. General Applicability

In order to assess the approach's general applicability to optical satellite ground imagery, we apply the trained model to 30-m GSD imagery from Landsat 8 and 1-m GSD imagery from satellite JL-1, as shown in Fig. 16. In both the cases, the trained SISR model increases the pixel resolution of the input



Fig. 16. Input images (top) and the results (bottom) of applying the SISR approach on (a) 30-m GSD imagery from Landsat 8 and (b) 1-m GSD imagery from satellite JL-1.

image, resulting in improved sharpness and clarity of features. The scale of the increase in resolution is inversely proportional to the resolution of the input image. While an increase in sharpness is evident in the 1-m GSD example shown in Fig. 16(b), the increase is slight compared to the increase seen with the 30-m GSD image from Landsat 8 shown in Fig. 16(a).

IV. DISCUSSION

A. Evolution of Small Satellite Capability

Constellations of small satellites provide reduced manufacturing costs, reduced launch costs, greater redundancy, greater resiliency against attack by adversaries, and a vastly improved ability to cover all areas of the globe from earth orbit. On the other hand, the reduced spacecraft size of small satellites results in a tradeoff in payload capacity relative to larger spacecraft with the same payload types and capabilities. By using ground-based methodologies once data products from small satellites are downlinked, small satellites can rival the capabilities of their larger cousins while maintaining the aforementioned benefits in costs, redundancy, and resiliency. Our capability demonstrated in this work dramatically increases the commercial and defense uses of small satellite optical ground imagery without any modifications to the spacecraft, payload, or mission operations procedures.

While only approximations can be made, a five times improvement in the pixel resolution of images from GEOSTare SV2 improves the GSD from the V4 telescope from 10.4 m to approximately 40 cm. This capability represents an industry advancement, allowing small satellite manufacturers and operators the ability to produce, launch, and operate single satellites

or constellations of small satellites that rival the capabilities of larger satellites. Moreover, the capability can also be integrated with larger satellites by using smaller, cheaper, and less power-hungry optical imaging payloads, freeing up physical space and resources for additional payloads or bus capabilities.

B. Capability Agnostic to Mission, Optical Payload, Data Source, or Spacecraft

Our developed capability has the ability to take any input image (tiled into square images 242×242 or less) and improve it to five times the original resolution to 1210×1210 pixels, irrespective of the mission, optical payload provider or type, the data source (mission operations or an archive of existing imagery), or spacecraft. The only requirements are that the images are: 1) three-channel RGB images and 2) 242×242 pixels or less when split into square tiles. In addition, while performance is lost as the size of the input image increases in size, our capability has the ability to take larger image tiles as inputs and provide high-resolution outputs up to the 1210×1210 pixel output image tile size. Developing a capability to apply across data sources, spacecraft, and optical imaging payloads dramatically increases the relevance of the capability and resulting product, as it can be applied across missions and across spacecraft, on both internal mission operations data sources and other sources such as NASA's Landsat optical imagery.

In addition, the developed capability can also be deployed through data reselling platforms for satellite imagery, improving the resolution of their existing and acquired image catalogues. The developed capability is a clear candidate for integration with spacecraft, payload, and mission offerings. That said, the developed SR capability also has significant utility through application on existing catalogues of satellite imagery of the earth surface, enhancing the defense, civil, and commercial readiness and value of available optical satellite imagery.

C. Achieving NIIRS Rating Level 4

As previously discussed, the NIIRS rating level 4 dictates that from aerial imagery "It is possible to identify farm buildings as barns, silos or residences; detect basketball or tennis courts in urban areas; identify individual tracks, rail pairs and control towers; detect jeep trails through grassland." As our goal is to achieve NIIRS rating level 4, we examine the results from our SISR approach against a set of images from both the V3 and V4 telescopes.

The V3 telescope provides a wide field of view with a 10.4-m GSD at 544 km in altitude. Fig. 17 shows a microscale ("zoomed-in") section of a GEOSTare SV2 image, which captures greenhouses among farmland in Saratov Oblast, Russia. The images in the figure show a pronounced improvement between the native-resolution GEOSTare SV2 image (center) and the improved resolution GEOSTare SV2 image (right). Notably, the cluster of greenhouses indicated in the figure were not discernible as single or multiple buildings in the original GEOSTare SV2 image, whereas the image produced with our SISR technique allows viewers to discern the structures as individual buildings. In addition, landscape textures and features are easily interpretable




Fig. 17. Google Earth (left) and improved resolution V3 telescope GEOSTare SV2 (right) optical satellite images of greenhouses amongst farmland in Saratov Oblast, Russia. The original native-resolution GEOSTare SV2 image is shown in the center.




Fig. 18. Google Earth (left) and improved resolution V4 telescope GEOSTare SV2 (right) optical satellite images of a tennis court adjacent to Beijing Capital Airport, China. The original native-resolution GEOSTare SV2 image is shown in the center.

in the SISR image, whereas they are difficult to interpret in the native-resolution original GEOSTare SV2 image.

The V4 telescope provides a narrow field of view with a 2-m GSD at 544 km in altitude. Fig. 18 shows a microscale (“zoomed-in”) section of a GEOSTare SV2 image, which captured a tennis court adjacent to Beijing Capital Airport in China. The figure indicates a dramatic improvement between the improved resolution GEOSTare SV2 image (on the right) versus the native-resolution image (in the center). Small buildings next to the tennis courts are clearer in the improved resolution image, and the borders of the courts’ playable area become more readily apparent as well. Nevertheless, the tennis courts are not easily identifiable without prior knowledge of the target area, possibly due to the lack of line and edge clarity or due to a lack of a color imager on GEOSTare SV2, which would add additional information for viewers of the resulting images. As such, while the improved resolution image shows a large improvement when compared with the native-resolution image, we believe we come close to—but not quite achieving—NIIRS rating level 4 with the V4 narrow-field-of-view telescope.

D. Future Work

There are several possible technical avenues to explore in order to achieve NIIRS rating level 4 from GEOSTare SV2 optical satellite imagery and our SISR technique. The first is to increase

the target resolution from five times to eight times or more. Currently, this is not possible with the batch sizes and available computing capabilities on AWS GovCloud, where the g4dn.12xlarge is the instance type with the largest amount of RAM available per GPU (16 GB). While the g5 instance types provide additional GPU computing capabilities and RAM, they are only available on AWS commercial services and not through GovCloud.

A second strategy is to utilize an alternative dataset instead of iSAID, where the optical images in the dataset are captured by satellites in orbit with a smaller GSD. A reduction in the GSD of the training data will result in greater detail in texture, shading, color, and other patterns in which to train a model, resulting in improved SISR performance. While the iSAID dataset provides satellite optical imagery at a GSD of 0.75 m with the JL-1 satellite, additional texture, edge, and pattern representations can be captured by our SISR approach using commercial satellite imagery with improved GSDs. In this respect, the target pixel resolution improvement of five times can be maintained while still improving the results of SISR.

Finally, additional SISR performance may be realized by utilizing alternative loss functions from the MSE (L2 loss). More recent work has indicated that gradient-based or gradient-guided loss functions may provide more desirable results in SISR applications [39], [40]. More specifically, SISR models may fail to recover sharp edges present in high-resolution imagery when leveraging MSE as the training loss function.

V. CONCLUSION

This work introduces the concept of leveraging SISR approaches to obtain high-resolution optical satellite imagery from low-cost satellite platforms frequently leveraged in commercial space. Traditionally, the cost of an optical telescope on board a satellite scales at roughly the square of the aperture, meaning that it costs four times as much to double the resolution of the imaging hardware. However, we show that SISR deep learning techniques can effectively improve the pixel resolution of an image by five times while correspondingly improving the NIIRS rating level, bringing low-cost satellites in closer competition with larger satellites with more capable optical payloads and without the expense of developing and constructing more capable telescopes. By providing low-cost small satellite platforms with the same capabilities as larger satellites, the cost for high-resolution in-orbit optical imagery is reduced alongside time to orbit (small satellites are built more quickly and typically launched more quickly through having a variety of launch vehicle options available). Methods are detailed, alongside results on Terran Orbital's GEOSTare SV2 satellite and satellites from other operators, such as the NASA. We report the performance of the results and detail areas of future work, which are likely to further improve the results of our described approach.

ACKNOWLEDGMENT

The authors would like to thank the mission operations team for GEOSTare SV2, Terran Orbital's mission operations and ground software teams, Terran Orbital's Information Technology and cloud engineering team, and Terran Orbital's business development team. The authors would also like to specifically thank Marc Bell, Wahid Azizpor, Marisa Jimenez, Richard Bays, Chris Shaffer, Mathias Koepfel, Akeem Valdez, and Chirag Patel.

REFERENCES

- [1] M. N. Sweeting, "Modern small satellites-changing the economics of space," *Proc. IEEE*, vol. 106, no. 3, pp. 343–361, Mar. 2018.
- [2] M. N. Sweeting, "The university of surrey UoSAT-2 spacecraft mission," *J. Inst. Electron. Radio Eng.*, vol. 57, no. 5S, 1987, Art. no. S99.
- [3] "Rivada space networks launches new low earth orbit constellation at satellite 2022. Rivada networks—Wireless unleashed," PR Newswire, Mar. 2022.
- [4] "Geostare SV2—Large satellite imaging quality, small satellite form factor and cost," Terran Orbital Corporation, Sep. 2022.
- [5] A. P. Cracknell and H. J. Kramer, "An overview of small satellites in remote sensing," *Int. J. Remote Sens.*, vol. 29, no. 15, pp. 4285–4337, Jul. 2008.
- [6] D. A. Holland, M. W. Thornton, and P. M. Atkinson, "Sub-pixel mapping of rural land cover objects from fine spatial resolution satellite sensor imagery using super-resolution pixel-swapping," *Int. J. Remote Sens.*, vol. 27, no. 3, pp. 473–491, Feb. 2006.
- [7] W. Shi et al., "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 1874–1883.
- [8] P. v. d. Smagt, C. Osendorfer, and H. Soyer, "Image super-resolution with fast approximate convolutional sparse coding," in *Proc. Int. Conf. Neural Inf. Process.*, 2014, pp. 250–257.
- [9] K. He, X. Tang, C. Dong, and C. Loy, "Image super-resolution using deep convolutional networks," 2015, *arXiv:1501.00092*.
- [10] J. Yang, W. Han, T. Huang, Z. Wang, and D. Liu, "Deep networks for image super-resolution with sparse prior," 2015, *arXiv:1507.08905*.
- [11] M. Irani, D. Glasner, and S. Bagon, "Super-resolution from a single image," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, 2009, pp. 349–356.
- [12] M. H. Yang, C. Y. Yang, and C. Ma, "Single-image super-resolution: A benchmark," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 372–386.
- [13] R. Fernandez-Beltran, J. Plaza, A. Plaza, J. Li, J. M. Haut, and M. E. Paoletti, "Remote sensing single-image superresolution based on a deep compendium model," *IEEE Geosci. Remote Sens. Lett.*, vol. 16, no. 9, pp. 1432–1436, Sep. 2019.
- [14] M. Shah and R. Gupta, "RescueNet: Joint building segmentation and damage assessment from satellite imagery," 2020, *arXiv:2004.07312*.
- [15] D. Hong et al., "More diverse means better: Multimodal deep learning meets remote-sensing imagery classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 5, pp. 4340–4354, May 2021.
- [16] S. Wei, T. Zhang, and S. Ji, "A concentric loop convolutional neural network for manual delineation-level building boundary segmentation from remote-sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–11, 2022.
- [17] S. Chen, Y. Ogawa, C. Zhao, and Y. Sekimoto, "Large-scale individual building extraction from open-source satellite imagery via super-resolution-based instance segmentation approach," *ISPRS J. Photogrammetry Remote Sens.*, vol. 195, pp. 129–152, Jan. 2023.
- [18] N. Buonaiuto et al., "Satellite identification imaging for small satellites using NVIDIA," in *Proc. Small Satell. Conf.*, 2017.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [20] T. Brox, O. Ronneberger, and P. Fischer, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervention*, 2015, pp. 234–241.
- [21] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248–255.
- [22] M. Everingham, L. V. Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, Sep. 2009.
- [23] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
- [24] A. Zisserman and K. Simonyan, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.
- [25] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, Jun., 2017.
- [26] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6517–6525.
- [27] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2980–2988.
- [28] L. C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 833–851.
- [29] P. Remagnino, M. M. D. Oghaz, and M. Razaak, "Enhanced single shot small object detector for aerial imagery using super-resolution, feature fusion and deconvolution," *Sensors*, vol. 22, no. 12, Jun. 2022, Art. no. 4339.
- [30] A. Gupta et al., "iSAID: A large-scale dataset for instance segmentation in aerial images," 2019, *arXiv:1905.12886*.
- [31] "The national imagery interpretability rating scale," *Nat. Collection Aerial Photography*, 2023.
- [32] G. van Rossum, "Python tutorial," Centrum Wiskunde and Informat., Amsterdam, The Netherlands, Tech. Rep. CS-R9526, May 1995.
- [33] S. van der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," *Comput. Sci. Eng.*, vol. 13, no. 2, pp. 22–30, 2011.
- [34] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 8026–8037.
- [35] J. Howard and S. Guggen, "FastAI: A layered API for deep learning," *Information*, vol. 11, no. 2, Feb. 2020, Art. no. 108.
- [36] J. D. Hunter, "Matplotlib: A 2D graphics environment," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 90–95, May/Jun., 2007.
- [37] Waskom et al., "Seaborn: Statistical data visualization," *Open J.*, vol. 6, no. 60, Jan. 2020, Art. no. 3021.
- [38] D. Hong, N. Yokoya, J. Chanussot, and X. X. Zhu, "An augmented linear mixing model to address spectral variability for hyperspectral unmixing," *IEEE Trans. Image Process.*, vol. 28, no. 4, pp. 1923–1938, Apr. 2019.

- [39] C. Ma, Y. Rao, Y. Cheng, C. Chen, J. Lu, and J. Zhou, "Structure-preserving super resolution with gradient guidance," 2020, *arXiv:2003.13081*.
- [40] L. Abrahamyan, A. M. Truong, W. Philips, and N. Deligiannis, "Gradient variance loss for structure-enhanced image super-resolution," 2022, *arXiv:2202.00997*.



Valentino Constantinou received the bachelor's of Science degree in economics from the University of Tennessee, Knoxville, TN, USA, in 2015, and the Masters of Science degree in analytics from Northwestern University, Evanston, IL, USA, in 2016.

He is an experienced Creative Problem Solver and Leader in analytics and artificial intelligence for the aerospace industry, serving as a Senior Data Scientist and Team Lead with Terran Orbital Corporation, Irvine, CA, USA. He was with the NASA Jet Propulsion Laboratory, Pasadena, CA, where he was the Principal Investigator to a multiyear alarm analytics effort, co-organized a monthly meetup of the lab's open-source developers (the Open Developer Meetup), and contributed to innovative research and technology development efforts. His efforts at the laboratory have resulted in numerous technology releases, publications, and patents. In 2017, he released as open source a deep-learning-based technology for scalable and accurate detection of anomalies from spacecraft, leveraging thousands of time-series telemetry streams—cited more than 900 times. In 2023, he authored or coauthored and released as open source a deep-learning-based technology for scalable and near-real-time detection of traveling ionospheric disturbances generated by gravity waves from tsunami waves. He is a regular contributor to open-source software in the areas of data science and machine learning and continues to conduct applied research and development of data science and machine and deep learning in aerospace.



Mark Hoffmann received the bachelor's degree in engineering and physics from Augustana College, Rock Island, IL, USA, in 2015, and the Masters of Science degree in analytics from North Carolina State University, Raleigh, NC, USA, in 2016.

He is an experienced Technologist in the areas of artificial intelligence and distributed systems. He has worked at the intersection of these areas in domains of intelligence, health care, aerospace, radio frequency (RF) signaling, and RecSys. Some notable works include several Defense Advanced Research Projects

Agency programs aimed at large-scale automated machine learning systems, optimization of NASA's Deep Space Network scheduling process, autonomous navigation of Martian rovers, and an online learning system that reverse engineers passive RF from electronics for identification. His current research interests include large-scale recommendation systems and pushing boundaries of Ads modeling by means of state-of-the-art artificial intelligence techniques.



Matthew Paterson was born in New England, USA. He received the first bachelor's degree in communication from the University of Maine, Orono, ME, USA, and the second bachelor's degree in English from the University of New Hampshire - Manchester, Manchester, NH, USA, both in 2007.

He is a Senior Data Scientist with Terran Orbital Corporation, Irvine, CA, USA, and an Artificial Intelligence Applications Engineer. He spent the last four years in consulting, working as a Data Scientist and Machine Learning Engineer with both small- and

enterprise-sized companies. He spent a year and a half teaching intro to Data Science courses. Since February 2021, he has been the Executive Director of The Spacelab, a space-focused 501(c)(3) nonprofit to help historically under-represented folks and others improve their careers in tech while discovering exoplanets and creating the Exoplanetarium.



Ali Mezher was born in Detroit, MI, USA. He received the Bachelor of Engineering degree in mechatronics engineering from Lebanese American University, Byblos, Lebanon, in 2022.

He is currently a Computer Vision and Robotics Engineer with Terran Orbital Corporation, Irvine, CA, USA. He was previously a Generative Artificial Intelligence and Computer Vision Researcher with the BMW Technology and Innovation Office, Munich, Germany, and the Partnerships Officer with the CodeBrave nonprofit organization. He authored or coauthored his first paper on "Visual Quality Assessment of Adversarially Attacked Images" and conducted and participated in numerous research projects relating to Mechanical Engineering Design and Simulation, Control Systems and Instrumentation, Robotics and Automation, and Biomechanics.



Brian Pak received the master's degree in statistics from San Diego State University, San Diego, CA, USA, in 2022.

With San Diego State University, he conducted research using convolutional neural networks to identify different stages of chronic pulmonary artery diseases. He is currently a Data Scientist with Terran Orbital Corporation, Irvine, CA, USA. His professional journey began with Kia, where he utilized machine learning and statistics to tackle safety issues in their vehicles. Before joining Kia, he contributed to Western Union by concentrating on the analysis of payment transactions to detect fraudulent activities.



Alexander Pertica received the bachelor's degree in physics from Cornell University, Ithaca, NY, USA, in 1984, and the master's degree in physics from the University of California, California, CA, USA, in 1986.

He is the Vice President for Space Domain Awareness Solutions with Terran Orbital Corporation, Irvine, CA, USA, developing small satellite solutions for intelligence, surveillance, and reconnaissance (ISR) and space situational awareness (SDA). Prior to joining Terran Orbital Corporation, he was

the Deputy Program Leader of Space Science and Security Program (SSSP) in Z-Program of the Lawrence Livermore National Laboratory (LLNL), Livermore, CA. The SSSP comprises a 50M+ USD portfolio of projects related to Space Science and Space National Security. As an Associate Program Leader in SSSP, he earlier managed execution of projects related to nanosatellite development for ISR and SDA. He led LLNL's efforts in the area of spectral signatures for missile defense warhead typing and kill assessment. He joined LLNL's Laser Program in 1986, working on laser materials for inertial confinement fusion. From 1990 to 2009, he was a Member of LLNL's Global Security Principal Directorate, working on the application of spectroscopic instrumentation to problems in the areas of missile defense, intelligence, and national security. Since 2012, he has been leading the Optical Sciences Group in Physics as a Group Leader. He was the Group Leader for Modeling, Simulation and Evaluation within the Advanced Technologies Division of the Global Security Principal Directorate. He was also the Project Leader for the Remote Optical Characterization Sensor Suite, an airborne remote sensor for missile defense test support. His research interests include space situational awareness, optical remote sensing, and optical and spectroscopic instrumentation.



Emily Milne received the Bachelor of Engineering degree in aerospace engineering from the Florida Institute of Technology, Melbourne, FL, USA, in 2022.

She was a Research Assistant with the Florida Institute of Technology, where she worked on software for a drone. She is currently a Spacecraft Operator with Terran Orbital Corporation, Irvine, CA, USA, where she focuses on the operation of GEOSTare SV2.