

# Synthetic Heterogeneous Anomaly and Maneuver - Neural Network Event Winnowing System (SHAM NNEWS)

**Dwight Temple**  
*ExoAnalytic Solutions*

## Abstract

With the immeasurable rise in space situational awareness data stream volume, elegant methods of high-flux analysis pale in comparison. Convolutional neural networks can support the variety of scenarios in addition to the sheer magnitude of incoming observations. However, hand-labeling these streams is infeasible. Using generative networks, semi-supervised learning methods are enabled and allow for exploration of not only known events in the dataset, but also can imagine and in-paint the sparse sections of the sample space. With this combination of technologies applied to the space domain, anomaly detection is enabled for well-known and unforeseen events. The debris generating event involving AMC-9 during July of 2017 shows the results of this technology stack applied to a pertinent and unseen vignette. This method, applied to ExoAnalytic Solution's increasingly voluminous data streams, enables real-time anomaly indications without a human in the loop.

## 1. Introduction

Neural networks, specifically convolutional neural networks (CNNs) and generative adversarial networks (GANs) [1], have permeated seemingly every scientific field. Automated vehicle navigation [2], medical scan anomaly detection [3], adversarial network attack/prevention [4][5][6], unknown state-space tracking [7], and recently space situational awareness [8] have shown provocative results for some of the most arduous applications. An industry standard is to utilize CNNs for supervised learning approaches for image classification and to use GANs for unsupervised learning where replication is the desired outcome. Proposed in this paper is a fusion of these methods to be operated in series to produce life-like representations of spacecraft anomalies in order to discriminate successfully these events from non-anomalous operating spacecraft behaviors. ExoAnalytic Solutions, with over 225 autonomous ground-based sensors, collects over 500,000 spacecraft images nightly; this vast database serves as an increasingly powerful resource bank that drives innovation for robust space-situational awareness (SSA) and neural network advancements. With recent debris-generating events occurring at geostationary orbit (GEO), proactive monitoring of spacecraft could permit avoidance of these disastrous events through specific indications and warnings (I&W). Specific events of interest include close approaches, glinting, plume generation, breakups and

other debris generating events.

## 2. Research Purpose

Discerning between nominal and anomalous operation with limited metrics, elevated noise, and star-obfuscated focal planes further complicates the problem of scene characterization. The hybrid architectures used and explored in this paper utilize disparate complete architectures as its components. For example, several GANs will be used for individual generative tasks, while stacked CNNs are explored for data segmentation and classification. Data utilized from ExoAnalytic Solutions will power a network of networks specifically operating on non-resolved satellite imagery. This paper describes the strengths of various neural networks for SSA applications and applies a set of CNNs and GANs to an example of a relevant application, focusing on the discrimination of close-approach maneuvering and nominal operation utilizing collected and synthetically generated data. The overarching research goal is to provide a novel framework capable of alerting SSA analysts and operates of anomalous behavior without a human-in-the-loop algorithm.


## 3. Background

### Convolutional Neural Network

Convolutions, a mathematically sound application in image and signal processing, can be applied in deep learning through learning of the filtering kernel. A CNN is composed of stacked layers of these learned kernels where the result is a classification label on an image. The depth of the stack determines the expressiveness and capacity of the image classifier by the complexity of the learned kernels. However, as the number of layers grows, the network becomes untrainable due to vanishing gradients. In 2012, the ImageNet architecture alleviated this issue using very

deep and stacked convolutions as opposed to feed forward layers; effectively, the number of gradient connections was reduced, and training was much more streamlined using newly optimized graphics processing unit (GPU) technology.

As a result, deep convolutional architectures have proliferated the literature in various forms. Utilizing residual connections allows for deeper and more expressive networks that can successfully discriminate images at a human-level capacity while still maintaining end-to-end trainability. In Fig. 1, examples of generic learned kernels can be observed.



**Fig. 1. Generic Convolutional Network Filters**

These trained kernels perform function specific actions within the neural network for object detection. For example, these filters may detect specific patterns such as vertical edges, rounded corners, color patterns, or background specifics relevant to the classification task at hand.

Generative Adversarial Network

In the same fashion, GANs also rapidly permeated industry as a novel method to sample from a latent space to generate fictitious, yet plausible, events. This property is coveted with regards to anomalous satellite behavior due to the limited sample size of captured elusive events. GANs function by staging two neural networks against each other in a minimax game. One network, the generator, creates samples from latent noise that it inputs into the other network, the discriminator, where the discriminator determines if the image is real or fake. Additionally, the discriminator is shown real images to also classify them as real or fictitious; the loss function of the generator is inversely proportional to the number of images the discriminator called fake, and the discriminator’s loss is the number of images it correctly classified as real and fake. This formula can be seen below in Equation 1.

$$\min_G \max_{D \in \mathcal{D}} \mathbb{E}_{x \sim \mathbb{P}_r} [D(x)] - \mathbb{E}_{\hat{x} \sim \mathbb{P}_g} [D(\hat{x})] \tag{1}$$

The GAN’s overall objective is to minimize the function with respect to the generator’s parameters and maximize the function with respect to the discriminator’s parameters.  $\hat{x}$  represents the generated output while  $x$  represents the real image inputs.

In Fig. 2 below, examples of generated imagery can be observed from a state-of-the-art GAN [17].




**Figure 2. Generated Results from PROGAN**

An objective of the methodology utilized is to use these generated plausible samples to further improve the classification capability of the CNN.

## 4. Supervised Data Collection


For empowering a CNN, massive amounts of labeled data are required for successful model convergence. However, labeled and standardized datasets of this nature are difficult to come by; for this reason, most deep learning applications are benchmarked against a select few premade sets (CIFAR10, ImageNet, MNIST). Standardization encompasses the image sizing, content, and labels applied to each sample. Ideally, each class of image would be equally represented within the corpus to equilibrate the gradient descent across each batch.

Regarding data volume, ExoAnalytic Solutions delivers over 500,000 unique satellite observations nightly, each of which a thumbnail of the resident space object (RSO) is stored. As a result, enormous databases are composed of an encompassing representation of events that occur at GEO. Unfortunately, sifting, sorting, and labeling this incoming data stream is infeasible and becomes increasingly difficult as operations expand. To manage this increasing complexity, semi-supervised learning can be utilized; this requires a baseline dataset at least. This initial task was generated using samples from one month of network observations from a sensor subset where the images were hand-labeled and sorted into two categories: anomalous and nominal. As mentioned, creating equal cross-sections of the event distribution is desirable. Anomalous events occurring at GEO are not commonplace, so maintaining the balance of classes required additional effort. To ameliorate this issue, anomalous samples were composed of observations of events with multiple RSOs, observations of the AMC-9 and Telkom-1 debris generating events [19], and observations with poor sensor measurements. Examples of anomalous images as inputs can be observed below in Fig. 3, and Fig. 4 depicts nominal examples utilized.



**Figure 3. Anomalous Collected Imagery**

Fig. 3 depicts a 12x48 grid of 32x32 (pixel) anomalous imagery collected from two sensors over the course of a month. Fig. 4 depicts the same dimension of images from the same sensors; however, these images are deemed nominal.



**Figure 4. Nominal Collected Imagery**

For the purposes of this application, pristine examples are of paramount importance because the network can only represent what it has been shown. Since the objective interest is to cleave as clean a separation as possible between anomalies and nominal behavior, the dataset must be highly representative of this divide. Situations where a partially occluded RSO, minor glints, or only slightly elevated signal-to-noise-ratio (SNR) were present were removed from the training corpus.

## 5. Semi-Supervised Classification and Dataset Augmentation

The primary objective is to efficiently and robustly categorize the incoming image stream as anomalous or nominal and report anomalous results to a specified satellite operator or analyst. This objective can be completed through the utilization of a semi-supervised deep learning approach to be further described.

Above, the distinctions between a supervised method such as a CNN and an unsupervised method such as a GAN were described. The former assists in classification and the latter in learning the unknown distribution of possible images. Because a converged and generalized GAN should produce images of sufficient quality to be deemed “real” by the network’s discriminator portion, the samples should be sufficient to augment the training dataset to the CNN for classification. As a result, one can freely sample from the latent space modeled by the GAN to generate life-like representations of possible images that have not actually occurred. As shown above in Fig. 2, the nonexistent celebrities generated are completely fictitious yet believable individuals. Therefore, the classification network can be exposed not only to instances of collected images, but also the potential outer boundaries of the unobserved sample space. Fig. 5 depicts the architecture framework used to accomplish this task.

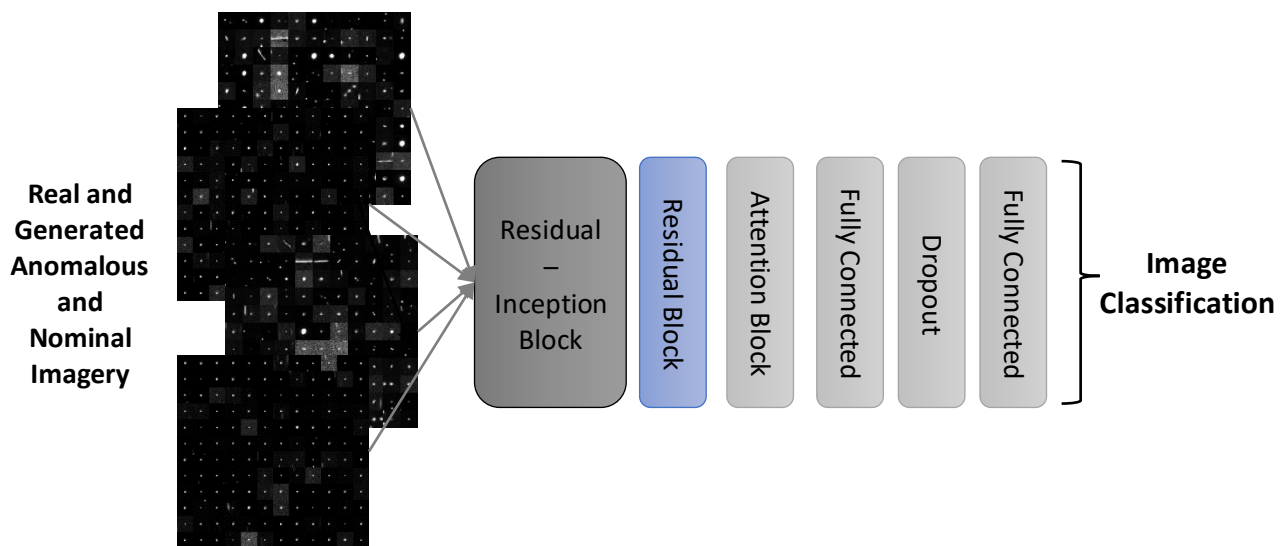



Figure 5. Convolutional Network Architecture

Further depiction of the specified blocks can be observed in Appendix I.

The system shown above is self-perpetuating once a stable and accurate classification model is completed. Beginning with the initial hand-labeled dataset, the converged GAN can then freely output representations to train a classifier. Most importantly, the converged classifier can be used to prescreen the unlabeled incoming image stream where the output will be a probability of anomalous or nominal. Based on this probability, high-confidence classifications can be labeled as their respective value, while low-confidence classifications can be screened via human-in-the-loop and labeled as their actual respective class. After this process is iterated over until a satisfactory classification rate is achieved, the result will be a classification network trained on a truth dataset numerous times larger and with more variance than the initial hand-labeled cases. Because of the perpetually increasing data volume, both the CNN and GAN will become robust and increasingly representative of the true state-space of the incoming images. Once a desired testing accuracy is met with the classification network on the holdout set, it can be used on a per-image basis with incoming imagery.

Fig. 6 depicts comparisons of GAN samples during the convergence process. Images on the left are from the initial training epochs while images on the right are from further epochs. Notably, the generator learns over iterations that anomalous images not only contain elevated SNR, but also additional objects and glints. GAN convergence was observed to only take roughly six hours on an NVIDIA 1060 6GB GPU; indeed, this model is lightweight and modular enough to be retrained daily for updated imagery databases.




**Figure 6. Convergence of Generative Network**

## 6. Analysis of Results

The networks described in previous sections were trained to converge their several processes of filtering inputs and outputs. The analysis of results completed below will depict both the training and testing classification accuracy and GAN samples. The testing set is composed of data *not* used during the training period and represents results from real-world operation on incoming data.

### GAN Results

While Fig. 6 depicts the convergence process of the GAN, Fig. 7 portrays sampling from the generator. These samples are completely fictitious representations of the “anomalous” latent manifold created by the GAN. Each row in the figure represents the smooth smearing between two potentially anomalous images. From left to right, the transition between variations can be observed.

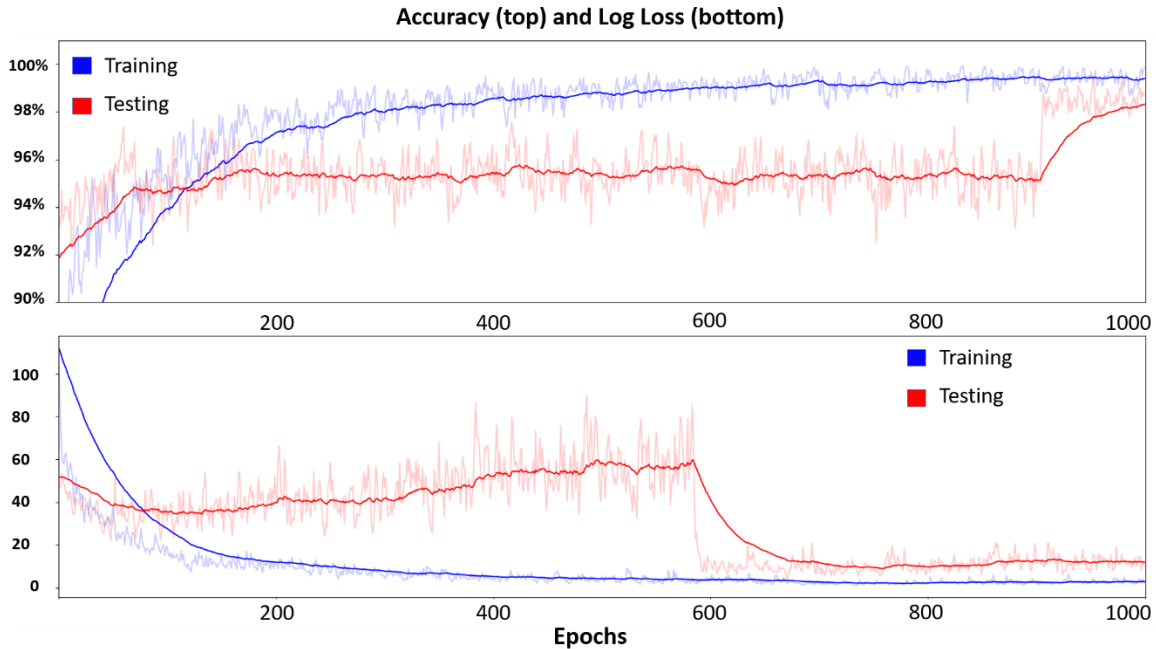


**Figure 7. Sampling from the Generator's Manifold**

Not only does the generator create anomalies similar to what it has been exposed to but also, it generates the unknown latent space between what it has previously observed. This power, harnessed properly, allows for a classification network to be constructed based upon the outer-bounds of what data has been collected.

## CNN Results

Fig. 8 depicts the comparisons of the overall dataset classification accuracy of the CNN trained on the dataset.



**Figure 8. CNN Convergence**

As seen in Fig. 8, the CNN converged smoothly as indicated by both the training and holdout testing losses.

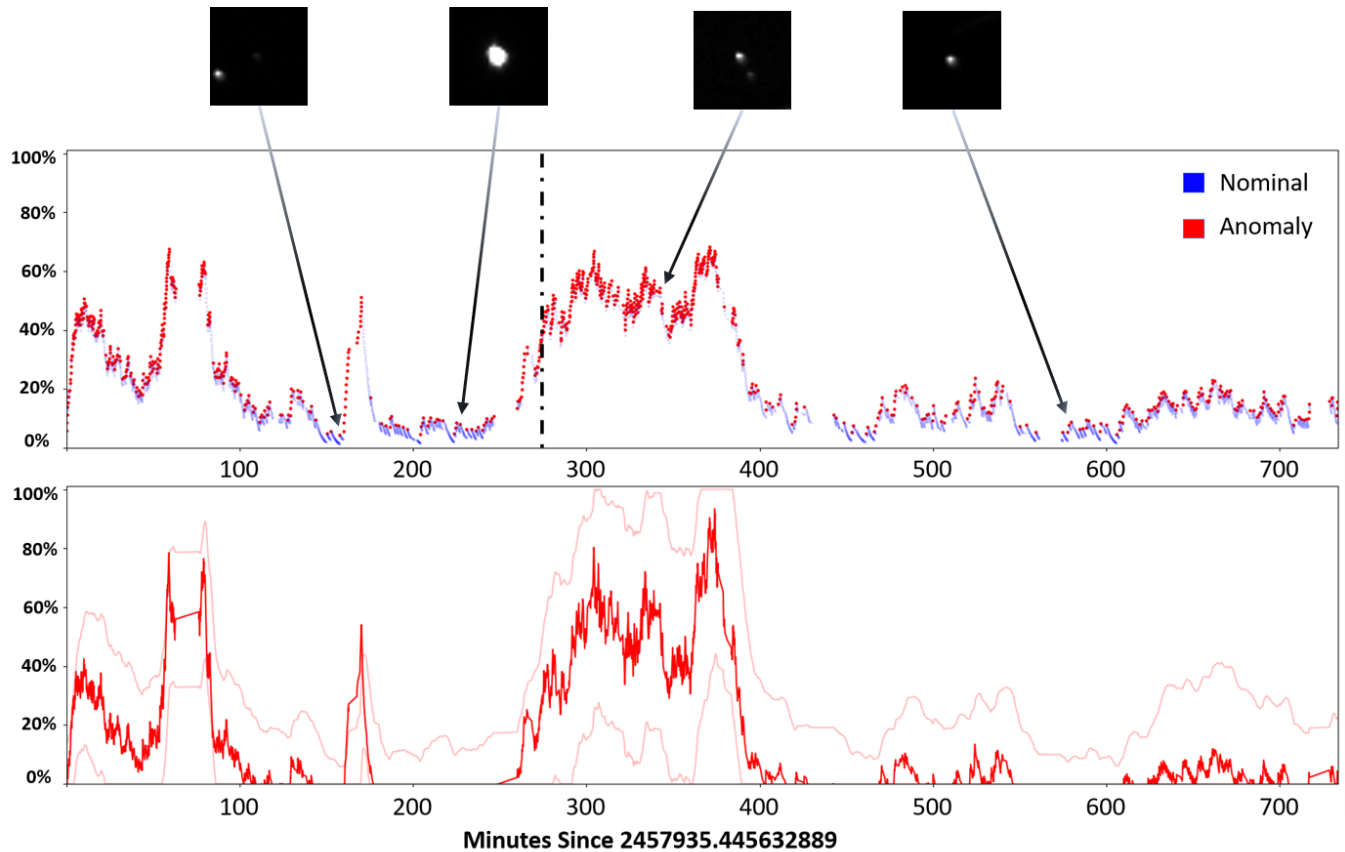
In testing the model, unlabeled images were used as inputs into the CNN; the outputs were the classification of the image and the probability that the image belongs to that respective class. A temporary graphic generated for visualization is included in Fig. 10 in Appendix I.

## 7. Implications and Discussion

While this combination of technologies may be an interesting technique for augmenting a training set, it must have practical implications in order to be a worthwhile pursuit. Fortunately, the application is readily available and can prove useful in numerous anomaly detection and I&W scenarios. In SSA, latency reduction is of paramount importance when making decisions regarding high-value assets such as communication satellites. Because the loop between the operator and observations can be time-consuming, actionable insights as soon as possible create vast differences in the decision-making process.

Tracking metric data alone can be an indicator of anomalous behavior by an RSO; however, if the metrics are within or close to the noise of the measurements taken, the anomaly may never surface to the operator or may surface far too late. Using photometric data coupled with a real-time anomaly detection suite allows for near instantaneous anomaly recognition and operator alert. To illustrate the utility in photometric anomaly detection, the AMC-9 event from June of 2017 will be analyzed using the CNN developed in this paper.

Metrically, the anomaly was not detected until a Julian date of  $\sim 2457935.638$  when a piece of debris was generated from the RSO. Utilizing the framework above, the anomaly was first detected concurrently and autonomously by specialized ExoAnalytic tools. With the anomaly detector running, ExoAnalytic could have known about it sooner and more effortlessly via detection of early indications. Potentially, the operator could have been alerted earlier or more observatories could have been tasked to observe the RSO. The timeline of observations of AMC-9 and their anomalous probabilities are shown in Fig. 9.



**Figure 9. Timeline of Anomalous Probability**

The black-dashed bar represents the time when the anomaly was first reported. As can be seen, there was another object in the scene detected ~ 150 minutes that persisted for about 20 minutes. Subsequently, the anomalous classifications began to spike as glinting was observed from AMC-9. The black bar is when heavy tumbling, glinting, and object fragmentation occurred. The correct classification of this event happened as well using the neural network. Finally, as the chaotic behavior of the RSO settled, anomalous calls were reserved to less frequent glints and debris around the RSO.

The main utility in automatic anomaly detection and alerts is the latency reduction *without* a human in the loop. Because the data flux is so tremendous for ExoAnalytic Solutions' telescope network, a non-leaky yet robust classifier must be able to prescreen information prior to it being sent to the analyst. The inclusion of this semi-supervised deep learning framework enhances the daily operations of SSA applications system-wide.

## 8. References

1. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. "Generative Adversarial Nets". In *NIPS*, 2014.
2. M. G. Bechtel, E. McElhiney, and H. Yun. "Deepicar: A low-cost deep neural network-based autonomous car." *arXiv preprint arXiv:1712.08644*, 2017.
3. Schlegl, T., Seeböck, P., Waldstein, S.M., Schmidt-Erfurth, U., Langs, G.: "Unsupervised anomaly detection with generative adversarial networks to guide marker discovery." *arXiv preprint, arXiv:1703.05921* 2017.
4. A. Kurakin and I. Goodfellow. "Defense against adversarial attack." *NIPS* 2017. <https://www.kaggle.com/c/nips2017-defense-against-adversarial-attack>, 2017.
5. A. Kurakin, I. Goodfellow, and S. Bengio. "Adversarial machine learning at scale." *arXiv preprint arXiv:1611.01236*, 2016.
6. A. Kurakin, I. Goodfellow, and S. Bengio. "Adversarial examples in the physical world." *arXiv preprint arXiv:1607.02533*, 2016.
7. M. Fraccaro, S. Kamronn, U. Paquet, and O. Winther. "A disentangled recognition and nonlinear dynamics model for unsupervised learning." *arXiv preprint arXiv:1710.05741*, 2017.
8. D. Temple, M. Poole, and M. Camp, "Network Enabled – Unresolved Residual Analysis and Learning", *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, held in Wailea, Maui, September 2017.
9. Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. "Deep Residual Learning for Image Recognition". *arXiv* 2015.
10. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. "Going deeper with convolutions." *CoRR*, abs/1409.4842, 2014.
11. Brachmann, A., and Redies, C. (2016). "Using convolutional neural network filters to measure left-right mirror symmetry in images." *Symmetry* 8:144. doi: 10.3390/sym8120144
12. M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN." *arXiv preprint arXiv:1701.07875*, 2017.
13. I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. "Improved training of wasserstein gans". In *NIPS*, 2017.
14. Zhang, H., Goodfellow, I., Metaxas, D. and Odena, A., 2018. "Self-Attention Generative Adversarial Networks." *arXiv preprint arXiv:1805.08318*.
15. O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. "ImageNet Large Scale Visual Recognition Challenge." *arXiv:1409.0575*, 2014.
16. A. Krizhevsky, I. Sutskever, and G. Hinton. "Imagenet classification with deep convolutional neural networks." In *NIPS*, 2012.
17. Martin Abadi, Ashish Agarwal, Paul Barham, Et al. "TensorFlow: Large-scale machine learning on heterogeneous systems", 2015.
18. T. Karras, T. Aila, S. Laine, and J. Lehtinen. "Progressive growing of gans for improved quality, stability, and variation." *arXiv preprint arXiv:1710.10196*, 2017.
19. Cunio, P. M., Bantel, M., Flewelling, B. R., Therien, W., Jeffries, Jr., M W., Montoya, M., Butler, R., and Hendrix, D. "Photometric and Other Analyses of Energetic Events Related to 2017 GEO RSO Anomalies." *AMOS*, 2017.



## 9. Appendix I

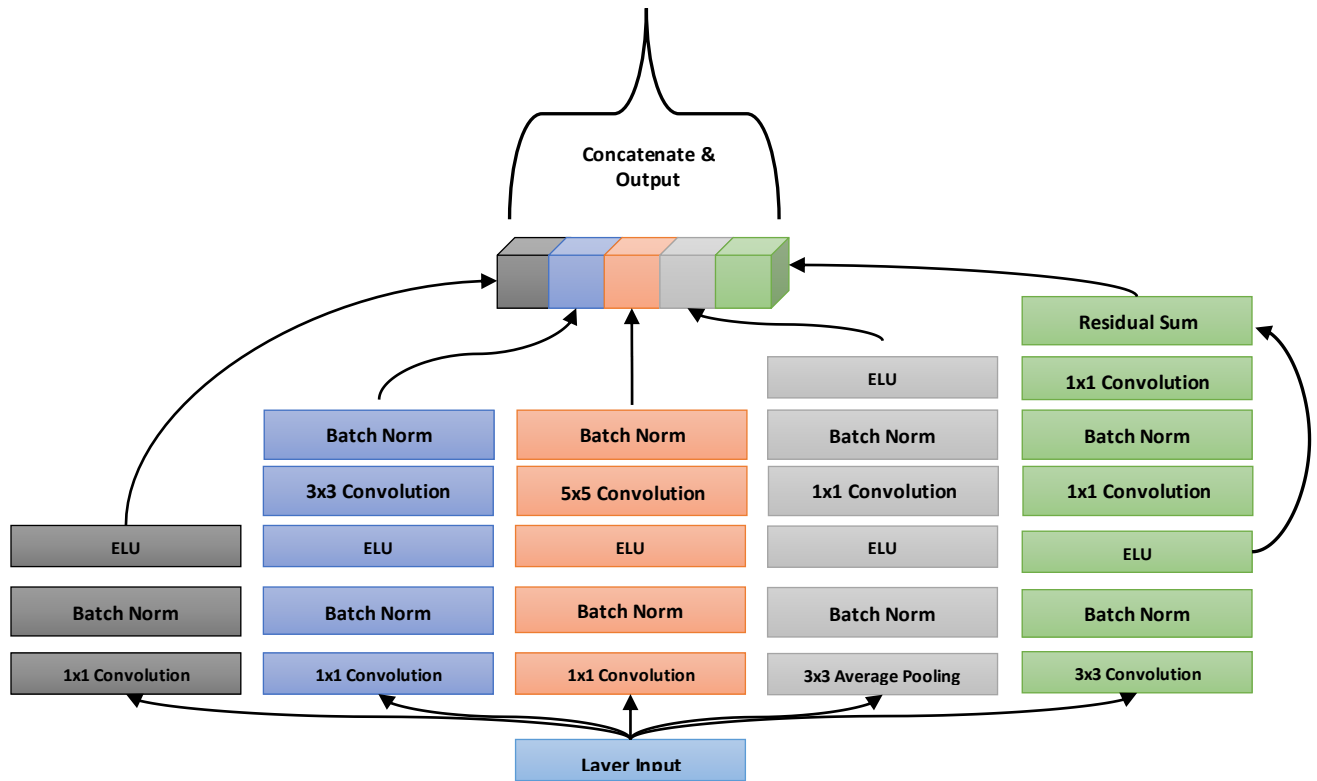


Figure 10. Residual Inception Block Component

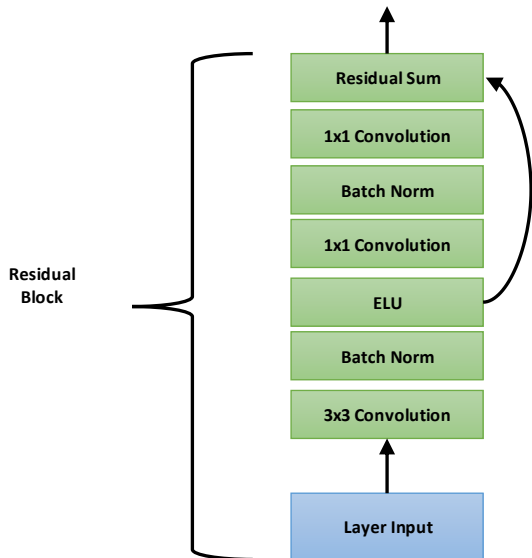


Figure 11. Residual Block Component

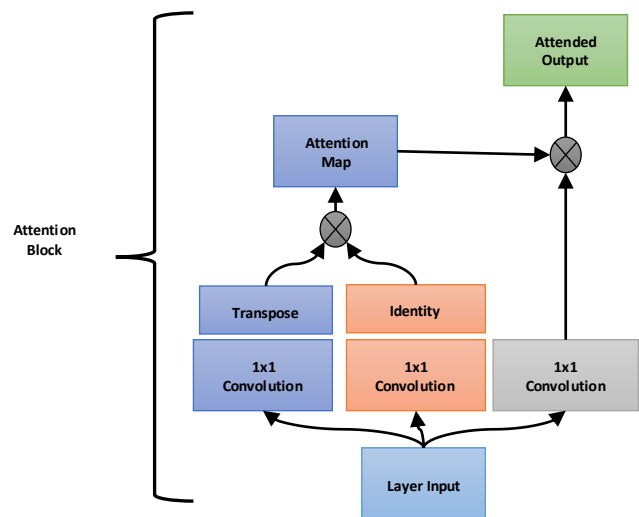


Figure 12. Attention Block Component