

**2020 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY  
SYMPOSIUM  
CROSS-CUTTING M&S CAPABILITIES TECHNICAL SESSION  
AUGUST 11-13, 2020 - NOVI, MICHIGAN**

**TRAINING ROBUST ANOMALY DETECTION USING ML-ENHANCED  
SIMULATIONS**

***Philip Feldman, PhD***

ASRC Federal, Beltsville, MD

***ABSTRACT***

*This paper describes the use of neural networks to enhance simulations for subsequent training of anomaly-detection systems. Simulations can provide edge conditions for anomaly detection which may be sparse or non-existent in real-world data. Simulations suffer, however, by producing data that is “too clean” resulting in anomaly detection systems that cannot transition from simulated data to actual conditions. Our approach enhances simulations using neural networks trained on real-world data to create outputs that are more realistic and variable than traditional simulations.*

**Citation:** P.Feldman, “Training robust anomaly detection using ML-Enhanced simulations”, In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium (GVSETS)*, NDIA, Novi, MI, Aug. 11-13, 2020.

**1. INTRODUCTION**

Creating autonomous vehicles that can perform optimally in unusual circumstances is a difficult problem in machine learning (ML). The reason for this is that most data is collected from systems functioning *normally*. For example, it is relatively straightforward to create a training set of typical rush hour traffic by simply equipping cars with cameras and driving them around in cities known to have traffic problems. But this is only a partial solution. In major evacuations, such as those for

hurricanes, traffic is often directed to use all available lanes. A self-driving car that is not trained for that possibility can be expected to behave in unpredictable ways. A neural network can easily “learn” to ignore such corner cases. For example, a network can be trained to drive one mile with perfect (99.9998%) accuracy if it assumes accidents simply do not happen [1].

This also happens with ML-based anomaly detection and response systems. The vast number of vehicles, from automobiles to satellites behave nominally for the vast majority of their functional

lifespan. Training for degraded modes requires vast amount of data being collected in a large number of malfunctioning states. Often, this data does not exist in sufficient quantity, and would be expensive to produce. One can imagine the paperwork required to slowly and rigorously destroy a collection of multi-million-dollar vehicles simply to train their diagnostic systems.

An effective solution to this problem is to use simulations [2], [3], [4]. Using synthetic data allows neural networks to be trained on edge cases in sufficient quantity such that the ML system can't develop undesirable biases. However, such simulations are often "too easy" for ML systems to understand, and fail in real world deployments [5].

To address this issue, we propose the use of machine learning to enhance the outputs of simple simulations, making them perform similarly to much more sophisticated simulators. In our current work with satellite anomaly detection for NASA and NOAA this technique is being developed to create realistic simulations for anomaly detection and classification, but we believe that it is broadly applicable. Briefly, the approach is as follows:

- 1) A simulator is constructed that approximately mimics the behavior of the target vehicle. This simulation can be quite coarse - for example a square wave can be used for nearly any periodic waveform, such the rotation of a wheel. This model does not have to include all systems on the target vehicle.
- 2) Data, either recorded from operational vehicles or from sophisticated, real-time simulators, is gathered in the course of normal operations. This data represents *baseline* behavior
- 3) The simple simulator is configured to generate its version of the baseline data, which is also recorded.
- 4) A first neural network is trained to *enhance* the simple data to match the general characteristics of the target data. This model learns to map the coarse behavior of the

simulator to a correct but generalized and unrealistically clean behavior. To add additional stochastic information to the output of this network, a second neural network is trained to replicate environmental contributions. The output of both neural networks are combined to produce a high-quality, realistic output.

Once trained, the enhanced simulator can infer realistic signals from a simulator that is running in a variety of "degraded" configurations. For example, shock absorbers can wear out. Air filters can become clogged. Subsystems can be crippled. Families of vehicles that are built on a common framework can be rapidly generated using the same simulator and different training data. This ability to quickly develop new capabilities that can be run rapidly on commodity hardware allows autonomous ML diagnostic systems to be trained effectively and at scale.

## 2. TELEMETRY EXAMPLE

An overview of the pipeline used to create lightweight, high-fidelity simulations is shown in Figure 1. For this example, the signals are synthesized sin waves with periods of 8 and 2 minutes. These signals are similar to those generated by rotating satellites (the shorter frequency) in orbit around another object (the longer frequency). To begin, we will generate an example signal and place it in our telemetry storage and retrieval system [6], shown as "Goal" in figure 1. Briefly, the steps involved in the process are:

- 1) Construct, lightweight, high-speed simulations
- 2) Generate approximate data
- 3) Train a neural network to map low-fidelity to high-fidelity data
- 4) High-fidelity simulation or real vehicle data source
- 5) Convert simulation to accurate, clean data
- 6) Environmental data source

- 7) Latent space data source
- 8) Train generative adversarial model (GAN) to create realistic environment influence
- 9) Generate environmental influences

The elements are then combined to produce the final, high-fidelity output (“Combined”).

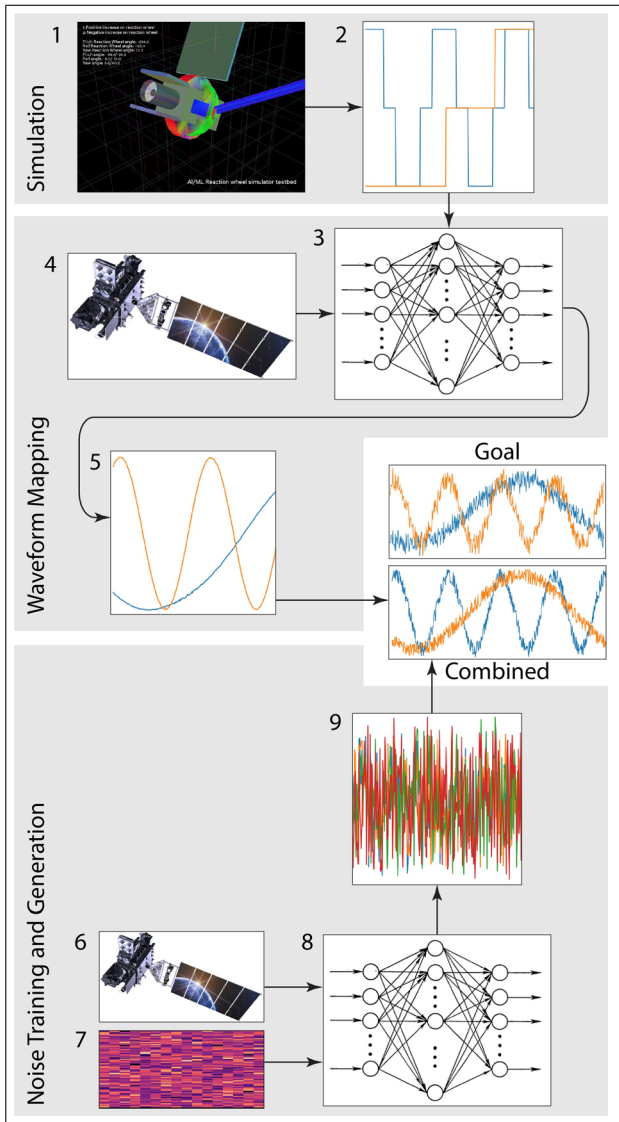


Figure 1: Enhanced simulation pipeline

We will apply the same pipeline and techniques to our example data that we would use for actual telemetry.

## 2.1. Simulation

For this development effort, we had access to highly sophisticated simulators for the NOAA GOES satellites and years of data. These simulators in many cases include the same software and often flight hardware. They are excellent for evaluating a particular set of options given a scenario and are extremely limited with respect to how much faster than real time they can operate.

ASRC Federal is in the process of developing simple software simulators that can be run in large numbers in the Cloud and much faster than real time. However, the cost for fast simulators is lower fidelity. So instead of the waveform shown as “Goal” in Figure 1 that would take 5-10 minutes to produce on a high-fidelity simulator, these simulators can generate the highly quantized data shown in Figure 2 (“Source”) in a few seconds.

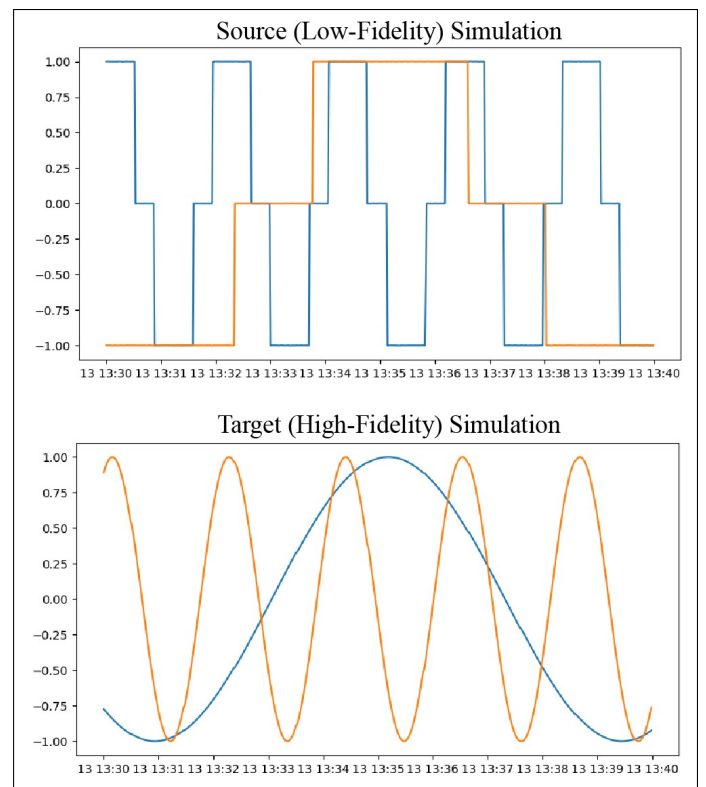


Figure 2: Training Source and Target Time Series

## 2.2. Waveform mapping

To transform the low fidelity output of the simulator into high-fidelity waveforms while maintaining the computational speed and memory footprint that simple simulations provide requires the training of two neural networks: The first is trained to map the output of the simple simulators to the output of the high-fidelity simulators or recorded activity from the actual satellite. All networks were implemented in Tensorflow version 2.1.0.

For this example, we developed a wide, shallow Multi-Layer Perceptron (MLP) network. The structure of the network is shown in Figure 3. It consists of four MLP layers (referred to in Tensorflow as “Dense”), with two inner, hidden layers (“Hidden\_1” and “Hidden\_2”) that then feed to an output layer. The size and number of dimensions are shown in the “Output Shape” column. Here we can see that these are one dimensional, with the number of neurons indicated by the second value in the tuple. These layers can be fed a variable number of input vectors, as specified by the first, “None” tag. The input and output layers are the size of the time series. The inner layers are wider, at 3,200 neurons. Wider networks are better at matching functions of this type [7]. The extra depth is required to match the multiple waveforms that the network has to learn. The last column, “Param #”, represents the total number of weights that the network will manipulate during the training sequence.

Layer (type)	Output Shape	Param #
Input (Dense)	(None, 500)	250500
Hidden_1 (Dense)	(None, 3200)	1603200
Hidden_2 (Dense)	(None, 3200)	10243200
Output (Dense)	(None, 500)	1600500

Figure 3: Enhancing MLP Neural Network

The model is trained by matching a large number of “source” time series such as those in Figure

2, with a corresponding set of high-fidelity “target” time series whose beginning and end are offset by a random amount so that all sample sizes are the same. After training the model for 40 epochs with a batch size of 15, we were able to produce the enhanced waveforms shown in Figure 4. These waveforms are produced by taking a specific time series of simulation data as a vector (Figure 2 “Source”) and mapping the input to an enhanced output vector (Figure 2 “Target”) of the same time. Timing for this output vector can be taken from the corresponding input vector element. Once trained, an input vector is multiplied by these weights to produce the enhanced values shown in figure 4.

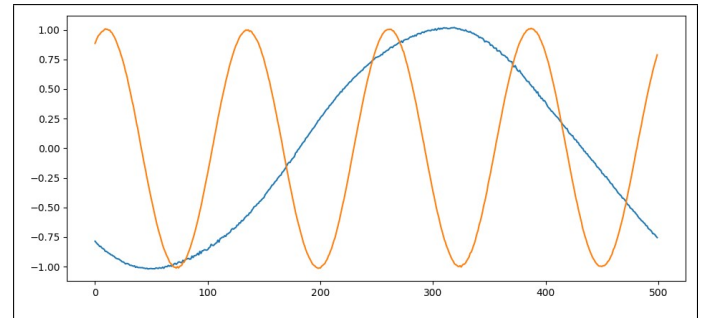


Figure 4: Enhanced Simulation

It is important to note that once the model is trained, that the inference that transforms the highly quantized simulation to the smooth, enhanced simulation is extremely fast, particularly when using hardware acceleration. This use of Neural Networks is what allows us to get high-fidelity results out of low-fidelity simulators without substantial speed or memory penalties.

However, we are now at the point that most high-fidelity simulation-based training systems encounter. The signal is too clean. An anomaly detection system trained on signals like these may not be able to discriminate between “normal” levels of noise and a genuine anomaly. It needs to be processed further to resemble the original waveforms in Figure 2.

### 2.3. Noise Training and Generation

In this approach, noise is trained independently using Generative Adversarial Neural Networks (GANs)

Generative adversarial learning is a technique where a *generative network* builds synthetic items (such as images) while the *discriminative network* attempts to distinguish the synthetic items from real ones take from a training set or *distribution* [8]. Typically, the generative network learns to map from a randomly generated latent space to the distribution of interest (such as pictures of faces), while the discriminative network tries to detect the synthetic items. The generative network’s training objective is to increase the error rate of the discriminative network by “fooling” the discriminator network through producing synthetic items that the discriminator thinks are real. This technique is quite capable of producing photorealistic results. The faces seen in Figures 5, 6, and 7, are *completely synthetic*, and were generated using the online StyleGAN2 generator [thispersondoesnotexist.com](http://thispersondoesnotexist.com) [9].



Figure 5



Figure 6



Figure 7

Noise and other stochastic environmental effects of real telemetry are extracted using a moving average filter [10]. This average is subtracted from the original signal, leaving the noise that needs to be simulated (Figure 8).

Our GAN for satellite telemetry is created by connecting the discriminator and generator networks together. The generator was built as shown in Figure 9. The 16-element latent vector is densely connected to a layer of 64 neurons. These are then normalized and spread to a layer of 500 neurons, the number of samples in the time series. This layer

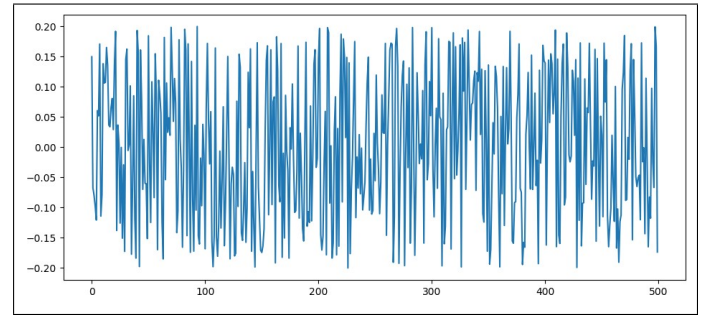


Figure 8: Extracted Source Noise

is reshaped to be compatible as the input layer of the discriminator.

Layer (type)	Output Shape	Param #
dense_2 (Dense)	(None, 64)	1088
batch_normalization (Batch Normalization)	(None, 64)	256
dense_3 (Dense)	(None, 32)	2080
reshape (Reshape)	(None, 32, 1)	0

Figure 9: Generator construction

The discriminator is provided with two sources that it must distinguish between:

- 1) A large number of real time series from the database whose beginning and end are offset by a random amount so that all sample sizes are the same.
- 2) A matching number of outputs from the generator, using the same format as the real input.

The discriminator (Figure 10) consists of one convolutional layer that merges input from 4 rows of inputs using a window 20 neurons wide with a stride of 4 neurons per step. This layer pools each convolution based on the maximum value of the sample. This configuration aids in finding patterns in signals. The layers are progressively narrowed to the final single neuron whose value determines if the input is considered real or fake. We use a

*binary crossentropy* loss function, which compares the summed errors across all classifications [11] combined with the *Adam* adaptive optimizer [12]. This value is then compared against the passed in values from the real and generated inputs to provide the information needed for training.

Layer (type)	Output Shape	Param #
conv1d (Conv1D)	(200, 4, 32)	672
dropout (Dropout)	(200, 4, 32)	0
global_max_pooling1d (Global (200, 32))		0
flatten (Flatten)	(200, 32)	0
dense (Dense)	(200, 25)	825
dropout_1 (Dropout)	(200, 25)	0
dense_1 (Dense)	(200, 1)	26

Figure 10: Discriminator construction

The real and generated data are tagged for discriminator training and evaluation. Real data has a tag of 1.0, while generated data has a tag of 0.0.

Training is divided into the training of the discriminator, and the training of the generator, as shown in Figure 11. In the first section, the discriminator as a standalone model is fed with equal amounts of data from the real data set and the generated data set. It then trains on the entire batch (200 rows) of real and generated data. After this pass, the discriminator’s weights are frozen, and the generator is trained as part of the entire GAN model. This allows the generator to be trained on the backpropagating error from the discriminator. To have the generator converge on realistic values, the tags for this pass are reversed, and the discriminator is “told” that the generated values are real. If it determines that they are *false*, then a distance is calculated that would adjust the weights towards the correct answer. Since the discriminator is frozen, the weights are only adjusted on the generator.

To match the qualities of this noise, our model

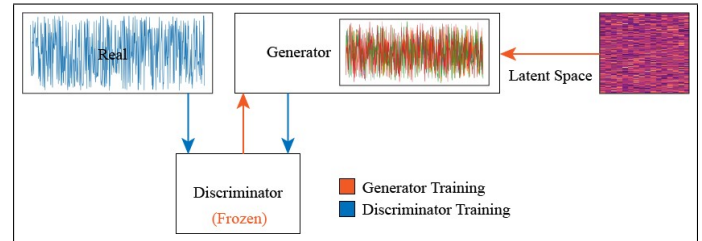


Figure 11: GAN Training

needed around 1,000 iterations. During this process, the *accuracy* – how many of the real and fake samples were correctly classified, and the *loss* – the normalized error across all classifications were sampled at 100-step intervals across the 1,000 iterations and are shown in Figure 12. It is important to remember when looking at this chart that the generator and the detector are engaged in an *adversarial* process, where the generator constantly tries to improve its ability to fool the detector, and the detector constantly tries to improve its ability to identify these forgeries. As we can see in the figure, the discriminator improves slightly faster than the generator, which is the goal of a GAN. If the two elements are too imbalanced, the system cannot learn effectively.

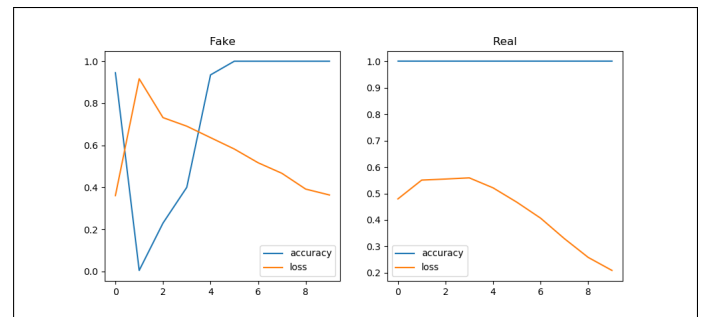


Figure 12: Accuracy/Loss for Real and Fake Classifications

After 1,000 iterations, the generated noise is sufficiently similar to the actual noise. The output of the simulator enhancing neural network can then be summed with the noise-generating neural network to produce the final signal, shown in Figure 13, below:

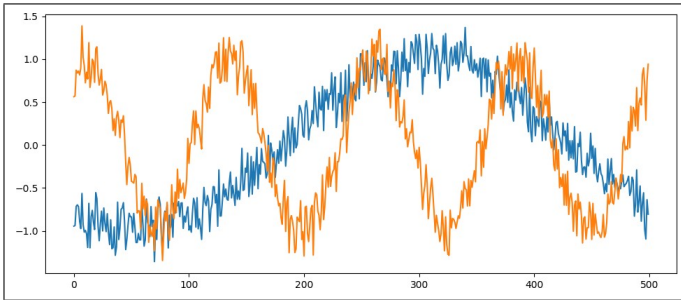


Figure 13: Simulator output, enhanced, with noise added

### 3. DISCUSSION

The majority of this particular research occurred during the peak of the COVID-19 NASA/NOAA response, and we were unable to access the high-fidelity simulators we were planning to use. As such the simulations described in this paper are based on simplified approximations. However, based on our experience modeling other satellite telemetry with small, rapidly-trained MLP networks [13] leads us to believe that the results here can be applied to actual telemetry data when it becomes available.

An issue that needs to be examined in more detail is the ability for the enhancing network to adapt low-fidelity signals to cases not covered by real telemetry or high-fidelity simulators. Because high-fidelity simulators are rare, they will only be used to explore likely problem spaces or respond to situations that occur on the actual vehicle. Unusual edge cases that show up quickly where only the low fidelity simulators are capable of responding will have to be researched more deeply to see if the data produced by the enhancing neural network is sufficiently valid.

An example of such an unusual scenario occurred in 2009, when US Airways flight 1549 struck a flock of geese shortly after takeoff resulting in a loss of power in both engines. With only a short period of time to evaluate potential options, the flight crew decided to ditch the plane in the Hudson river, saving all passengers [14].

What would have happened if the crew had

been less experienced? Could there be a way to evaluate options for these types of cases where time is critical and experience limited or nonexistent? If enhanced simulators can be built to be small and fast enough to run at many times normal speed and in parallel, it may be possible to automate a response to a Mayday request by starting an always-available cluster of simulations to evaluate potential best options given edge-case degraded modes. In essence, multiple reinforcement learning simulations are set up with the objective function being in the case of Flight 1549, a safe landing.

Such simulations need not be limited to satellites or civil aviation. Simulation and prediction of degraded ground vehicle behavior ranges from situations as specific as overheating train axles [15] to predicting traffic [16]. Combat often involves ground vehicles operating individually or in groups in degraded modes that cannot be anticipated. An approach to adapting quickly to these unanticipated situations using scalable high-fidelity simulation may make for a more adaptive combat capability that is able to adjust to changing conditions faster than the Adversary.

### 4. CONCLUSIONS

All machine learning depends on large volumes of data. Creating a pipeline for providing synthetic data on demand is a market that is currently worth tens of millions of \$US annually that is likely to only increase over time. ASRC is developing systems to provide synthetic data at scale. Synthetic data allows organizations to be independent of data sources with potential limitations and foreign complications.

Machine learning models are useless without data, and diverse data can make the same model applicable in diverse contexts. Even if progress ceased in the development of more sophisticated models, machine learning could be effectively applied to new domains simply by training current state-of-the-art models with new, well understood and balanced datasets.

Simulation as a way of creating usable assets is currently being done in an ad-hoc basis in the AI/ML community. Particularly for the government user, it is often the only secure way to generate the amounts of data needed for the effective training of unusual models, such as satellites. In this paper, we have shown that it may be feasible to produce large amounts of simulated data that can in turn be used to train machine learning systems to recognize and adapt to rare and unlikely situations. Future work will focus on increasing the range, scale, and sophistication of these types of simulations.

## REFERENCES

- [1] NHTSA National Center for Statistics and Analysis, “2018 fatal motor vehicle crashes: Overview,” 2018.
- [2] C. J. Pretorius, M. C. du Plessis, and C. B. Cilliers, “Simulating robots without conventional physics: A neural network approach,” *Journal of Intelligent & Robotic Systems*, vol. 71, no. 3-4, pp. 319–348, 2013.
- [3] D. Zadok, T. Hirshberg, A. Biran, K. Radinsky, and A. Kapoor, “Explorations and lessons learned in building an autonomous formula SAE car from simulations,” *arXiv preprint arXiv:1905.05940*, 2019.
- [4] A. Bewley, J. Rigley, Y. Liu, J. Hawke, R. Shen, V.-D. Lam, and A. Kendall, “Learning to drive from simulation without real world labels,” 2018.
- [5] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, “Sim-to-real: Learning agile locomotion for quadruped robots,” 2018.
- [6] Staff, “Influxdb 2.0: Complete time series toolkit — influxdata,” [Accessed 12-June-2020]. [Online]. Available: <https://www.influxdata.com/products/influxdb-overview/influxdb-2-0/>
- [7] S. Zagoruyko and N. Komodakis, “Wide residual networks,” 2016. [Online]. Available: <https://arxiv.org/abs/1605.07146>
- [8] Z. Wang, J. Chen, and S. C. Hoi, “Deep learning for image super-resolution: A survey,” *arXiv preprint arXiv:1902.06068*, 2019.
- [9] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410.
- [10] Staff, “pandas.core.window.rolling.rolling.mean,” 2020. [Online]. Available: <https://pandas.pydata.org/docs/reference/api/pandas.core.window.rolling.Rolling.mean.html>
- [11] R. Y. Rubinstein and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media, 2013.
- [12] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [13] Z. Li, “Machine learning in spacecraft ground systems,” in *2017 6th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE, 2017, pp. 76–81.
- [14] Wikipedia contributors, “Us airways flight 1549 — Wikipedia, the free encyclopedia,” 2020, [accessed 12-June-2020]. [Online]. Available: [https://en.wikipedia.org/w/index.php?title=US\\_Airways\\_Flight\\_1549&oldid=961609724](https://en.wikipedia.org/w/index.php?title=US_Airways_Flight_1549&oldid=961609724)
- [15] X. Yang, H. Dong, J. Man, F. Chen, L. Zhen, L. Jia, and Y. Qin, “Research on temperature prediction for axles of rail vehicle based on LSTM,” in *International Conference on Electrical and Information Technologies for Rail Transportation*. Springer, 2019, pp. 685–696.
- [16] M. Jiber, I. Lamouik, Y. Ali, and M. A. Sabri, “Traffic flow prediction using neural network,” in *2018 International Conference on Intelligent Systems and Computer Vision (ISCV)*. IEEE, 2018, pp. 1–4.