2023 NDIA MICHIGAN CHAPTER
GROUND VEHICLE SYSTEMS ENGINEERING
AND TECHNOLOGY SYMPOSIUM
MODELING SIMULATION AND SOFTWARE (MS2) TECHNICAL SESSION
AUGUST 15-17, 2023 - NOVI, MICHIGAN

SCENE ADJUSTMENT TOOLS FOR AUTONOMY SIMULATION

Joseph Auchter, Robert Brothers, Daniel Beer

Southwest Research Institute, San Antonio, TX

ABSTRACT

This paper introduces a set of Simulation Scene Adjustment Tools which allow users to adjust relevant geo-environmental parameters to create an unlimited number of new scenes that are based on, and share relevant characteristics with, the original. We also introduce a set of novel quantitative metrics that compare the simulation scenes from the perspective of an automated vehicle system. These tools are useful for automated vehicle development and testing to alleviate the problem of having only a few simulation scenes with which to prove out the system.

Citation: J. Auchter, R. Brothers, D. Beer, "Scene Adjustment Tools for Autonomy Simulation," In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium* (GVSETS), NDIA, Novi, MI, Aug. 15-17, 2023.

1. INTRODUCTION

Simulation tools are useful for development and validation of automated vehicle (AV) systems. To thoroughly test new algorithms and techniques in simulation, researchers require numerous test scenes that share certain relevant characteristics with a given real (geo-specific) location where the vehicles will be operating. Having such a variety avoids the pitfall of tuning and training a system that works very well in one scene but fails in a scene that differs from the original. A need exists for novel simulation scene generation tools to alleviate this over-training problem for AV systems [1].

Recent advancements in simulation systems for off-road AV development and testing include the ability to randomly generate scenes with realistic-looking terrain elevation and foliage distributions that can be configured by the user [2]. The approach we describe in this work is unique because our tools allow for perturbations (or "adjustments") to be applied to original geographic information systems (GIS) data that can represent a real-world location of interest. 3D simulation scenes can then be created from the original and perturbed GIS data to give the user a variety of scenes which share similarities with, but are also different from, the original scene.

2. PROJECT OBJECTIVE AND TECHNICAL APPROACH

Southwest Research Institute (SwRI) undertook an internal research project to develop scene generation and adjustment tools to aid in AV development and testing. The project had three objectives:

- Create a system that rapidly generates 3D simulation scenes from GIS data.

This work was funded by a Southwest Research Institute Internal Research Project

- Develop novel tools to adjust, or perturb, relevant geo-environmental variables to create one or more synthetic scenes that are similar (from the perspective of the AV) to a given geo-specific scene but are different enough to exercise the AV system in ways the original scene does not.
- Establish novel quantitative metrics that compare the simulation scenes to each other from the perspective of the AV.

Our technical approach to achieve these objectives is discussed in the following sections.

1.1. GIS Data to 3D Simulation Scene Pipeline

For this work we used the Unreal Engine 4 (UE4) [3] simulation engine. Unreal Engine was chosen because it is widely used in the simulation industry due to its realistic visuals, tools for creating scenes, and plugin architecture for connecting with other products. This section describes the procedure used to create 3D scenes in UE4 from original GIS data.

To evaluate the scene adjustment algorithms and scene comparison metrics for this research, we created a 3D scene of the SwRI Off-Road Test Facility at our San Antonio, Texas, headquarters. We used three sources of GIS data to build the scene:

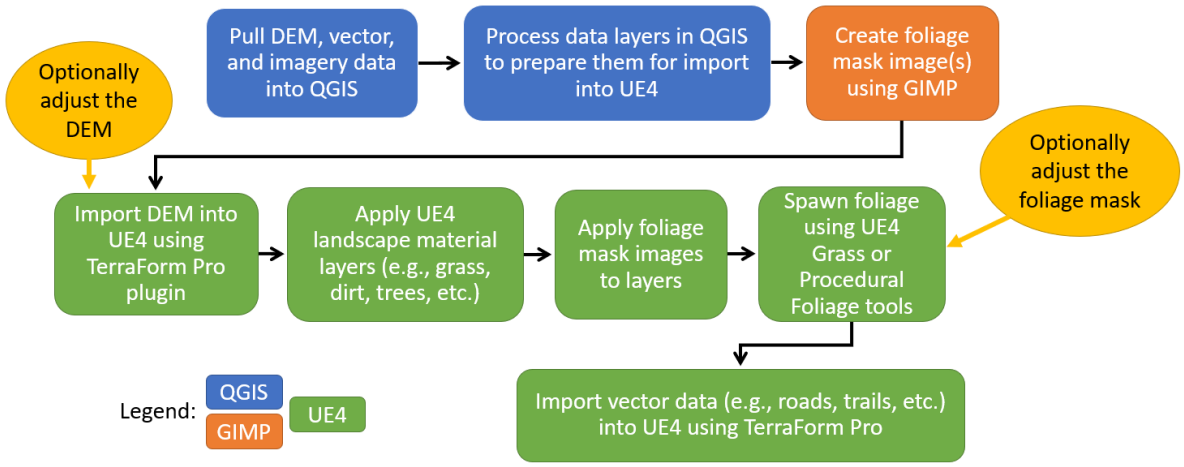

Figure 1: Flow chart for creating the 3D scene from raw GIS data.

- Digital Elevation Models (DEMs) from high-resolution aerial scans made available by the San Antonio River Authority [4].
- Vector data for roads, trails, and other similar features, available from OpenStreetMap [5].
- Color satellite imagery, available from the U.S. Geological Survey [6].

We used QGIS [7] to gather and process the GIS data. To create foliage mask images from the satellite imagery, we used GNU Image Manipulation Program (GIMP) [8]. Since UE4 does not natively handle GIS data files such as GeoTIFF (an image file format with geo-referencing metadata) files and Esri ShapeFiles [9], we used a commercial UE4 plugin called TerraForm Pro [10] to import these GIS files to create the terrain and place features like foliage and trails in geographically accurate locations.

The flowchart in Figure 1 describes the process to create the simulated 3D scene from the original GIS data, including optional adjustments to geo-environmental elements

(terrain elevation and foliage placement). An example of a portion of the SwRI Off Road Testing Facility map in QGIS and the corresponding section of the 3D scene with accurate terrain elevation, foliage distribution, and trail locations is shown in Figure 2.

Figure 2: Top: a multi-layer GIS map of a portion of the SwRI Off-Road Test Facility assembled in QGIS. Bottom: an overhead view of the matching portion of the completed 3D scene in UE4.

1.2. Algorithms to Adjust GIS Data

One of the major objectives of this research was to create and evaluate new ways to adjust real-world GIS data in user-configurable ways to create an arbitrary number of synthetic scenes that share similarities with the original scene. We created a set of "Scene Adjustment Tools" that allow users to adjust two aspects of the scene that are of high importance to off-road AV performance: terrain elevation and foliage placement.

First, we discuss our algorithm for adjusting the terrain elevation. Procedural map and texture generation from pseudo-random noise has been popular in the game development community for many years [11]. The task of game map generation is very similar to the goal of adjusting an existing simulation map: noise must be added to a base elevation map layer in a way that appears like naturally occurring terrain features. White noise, or noise that has a Gaussian distribution, is often a poor choice for mimicking a signal found in nature. "Pink" or "red" color noise is a more appropriate choice for modeling natural features like clouds or terrain. Pink noise is characterized by equal energy in each octave (power spectral density of 1/f), while red noise is characterized by having more power in the lower part of the frequency spectrum (power spectral density of $1/f^e$, $e > 1$).

Our approach to adjusting DEMs was to generate noise in a way that could be used as an entirely new digital elevation map and then combine the noise with the original DEM using a weighted blending. We first explored using a simplex noise generator to create these additive disturbances. Generating simplex noise is relatively fast for small streams of random numbers; however, this process does not scale well to the size of the high-resolution (~5 megapixel) digital elevation maps used in this research.

We next attempted to generate the same type of noise using a technique called frequency synthesis, which was computationally time-efficient for high-resolution images. Frequency synthesis uses the knowledge of the "shape" of the frequency response of the desired output noise to filter an easy-to-generate input noise source [12]. For our pink noise synthesis algorithm, we first generated a white noise image of the desired resolution using a Gaussian noise generator. See Figure 3 for an example Gaussian noise image.

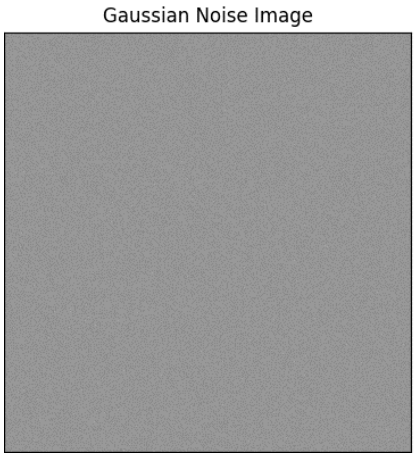

Figure 3: An example white noise image.

We then produced the Discrete Fourier Transform (DFT) of the white noise image using optimized implementations of the DFT algorithm. A simple shaping filter was applied to the DFT image by multiplying each pixel by $1/\left(freq_{pixel}{}^{exp_{shaping}}\right)$, where $freq_{pixel}$ is the pixel value in the DFT image, and $exp_{shaping}$ allows us to vary the color of the output noise (2.0 for red noise and 1.0 for pink noise).

After the shaping filter was applied, the DFT image could be converted back to the spatial domain using an Inverse DFT (IDFT), which resulted in a noise image with the desired frequency characteristics, such as the one shown in Figure 4.

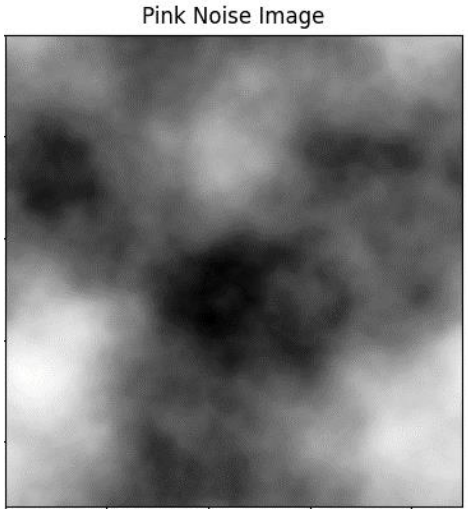

Figure 4: A pink noise image in the spatial domain.

By generating the original white noise image with zero mean value, we generated constructive and destructive adjustments when the noise image was added to the original DEM. The variance of the original white noise generation can be scaled to fit the desired goals for adjusting the DEM. For example, for a small-scale robotic platform, a reasonable elevation adjustment might be a maximum of ±5 meters. The variance can be calculated accordingly.

Directly adding the pink noise to the input DEM sometimes resulted in unnatural features in the output DEM, which are not useful for developing and evaluating autonomy systems. We experimented with different levels of weighted blending of the noise and original DEMs to preserve some of the original's sharp features or eliminate roughness in certain areas, for example. An appropriate level of blending depends on the original or input DEM and the desired level of noise in the output, so we made the blending configurable in our DEM adjustment tool.

Our foliage mask adjustments were done using different techniques. "Blue" noise can be used to generate assets like trees that are sparse or scattered through a scene. Blue noise is characterized by containing more signal power in the higher portion of the frequency spectrum. Poisson disk sampling is a simple algorithm for generating realistic blue noise and is popular in game development [13]. Figure 5 is blue noise generated from the classic Poisson disk sampling algorithm that uses a uniform distance sampling in each iteration.

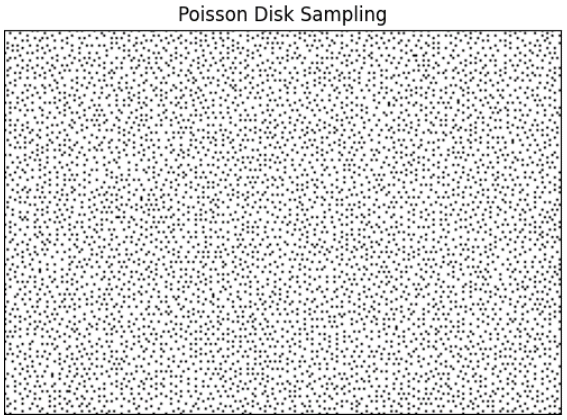

Figure 5: An example image produced by Poisson disk sampling.

In our 3D scene generation pipeline (Figure 1), we used mask images that defined the placement of trees, tall grass, or other terrain objects. We chose to adjust the placement and layout of these objects by manipulating these mask images. Our goal was to generate new mask images that maintained large clumps of objects (low-frequency components) but also contained some high-frequency dispersion of new objects.

A useful, novel modification to the classic Poisson disk sampling was developed during this project to help preserve the desired low-frequency characteristics of the original mask image while realistically modifying the high-frequency components. This modification used an existing image's intensity values to vary the sampling distance between a minimum and maximum value for corresponding pixels of the existing image and the output image the algorithm generates. In practical terms, the modified algorithm uses the existing trees and tall grass mask images to drive the blue noise generation. The results give us generated outputs with similar clumps of objects in areas of high density and sparse blue noise points in areas of low density in the original image. Figure 6 shows an original tree placement mask image and an adjusted version of it.

Using sparse input images sometimes caused the algorithm to fail to find any valid points to generate, which ended the algorithm early and resulted in a nearly empty generated image. To avoid this null result, we added a minimum sampling percentage parameter to reseed the algorithm if the output contains fewer samples than a desired number of pixels from the input image. This minimum sampling percentage parameter can be tuned to generate reasonable output

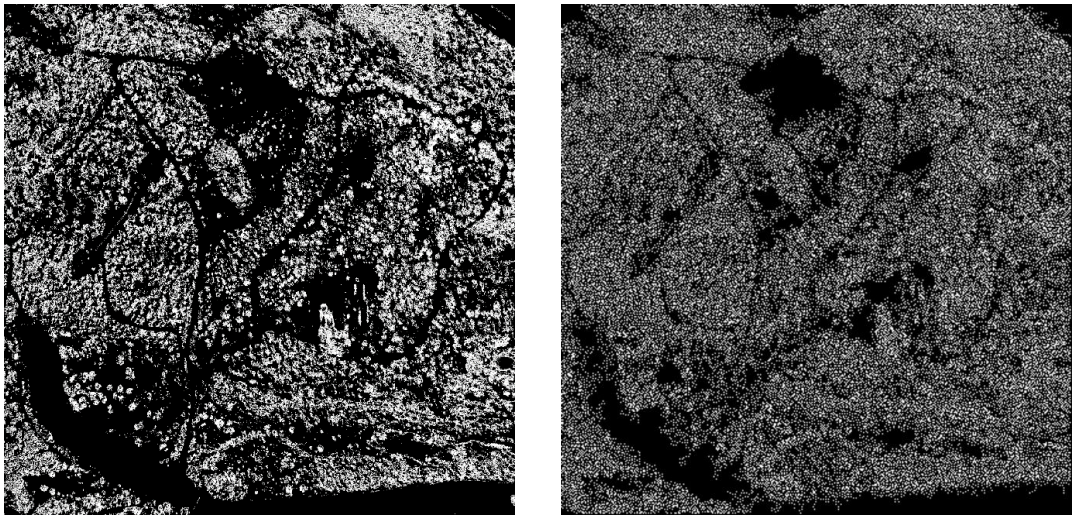

Figure 6: Left: An original tree mask image. Right: adjusted version. White areas indicate where trees will be placed in the 3D scene.

images from any level of sparsity in the input images.

1.3. Metrics to Compare Scenes

An important aspect of generating numerous simulation scenes using the tools described above is the ability to compare the scenes to one another quantitatively, from the perspective of the AV. This comparison is useful because a user might want to test the AV system in, say, 100 scenes that share similar characteristics to a real location in which the AV will operate. Quantitative measures are needed to assess how similar the scenes are to the original to avoid using scenes so similar as to provoke identical responses or so different as to be irrelevant. This section discusses our development of several novel scene comparison metrics.

A costmap is a top-down view of a geographic area in the vicinity of the AV. The area is divided into grid cells, and each cell is assigned a numerical value indicating some metric about the location, such as the cost of the cell as determined by whether the AV can expect to be able to traverse it. The costs can be generated from a variety of fused, processed sensor data such as lidar returns and electro-optical (camera) images. Higher numerical values indicate higher cost and therefore more risk to the AV. The costmap is of crucial importance to AV behavior, as the AV makes many of its path planning decisions based on it. Therefore, the costmap is an appropriate place to make comparisons about how the AV perceives its environment. This decision also avoids the need to tune our algorithms for specific sensor payloads.

Figure 7 shows an example costmap based on the first- and second-order gradients of the DEM of a portion of SwRI's Off-Road Test Facility. The first-order gradient is the local slope of the terrain, and the second-order gradient is local terrain roughness. This costmap was used to calculate the metrics described below.

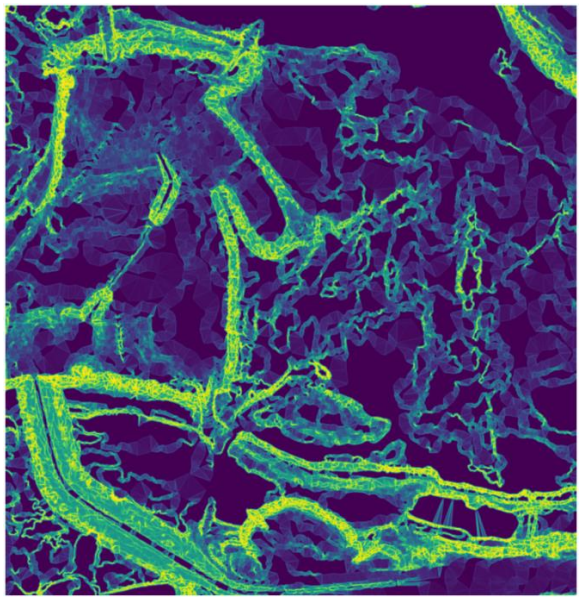

Figure 7: Example costmap of the SwRI Off Road Test Facility. Dark areas are low cost, and bright yellow areas are high cost.

First, we divided the costs of individual costmap cells into histogram bins (for example, in cost categories such as low, medium, high, and fatal). Histograms of these bins for two different scenes might look something like those shown in Figure 8.

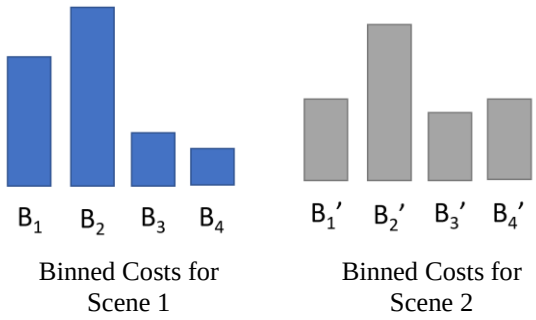

Figure 8: An example of costs that have been binned into four categories for two different scenes.

To assess how similar two scenes are to each other as perceived by the AV, we compared the costmaps associated with each scene using three quantitative metrics:

1. **Traversability Ratio Metric.** Bin the costs and calculate the ratios between the corresponding bins from one scene to

another, as shown in Equation (1). If the ratio is close to 1, this indicates that the AV perceives the two scenes to be similar.

$$R_1 = \frac{B_{1_{scene\ 2}}}{B_{1_{scene\ 1}}} = \frac{B_1'}{B_1} \qquad (1)$$

2. **Blob Size Metric.** For each cost bin in a costmap image, perform connected component labeling to determine the area (in pixels) of each blob. Calculate the average blob size over the whole costmap image. Compare these averages among the different scenes. Closer averages indicate that the AV perceives the scenes to be similar. Figure 9 shows a flow chart of this calculation and comparison process.

3. **L_2 Metric.** Bin the costs of the costmaps and create histograms. Compare the histograms from different scenes to each other using the L_2 distance (also known as the Euclidean distance). The distance is calculated as in Equation (2), where *n* is the total number of bins in each histogram. Smaller distances indicate that the AV perceives the two scenes to be similar.

$$d(B, B') = \sqrt{\sum_{i=1}^{n} (B_i' - B_i)^2} \qquad (2)$$

3. RESULTS

Figure 10 shows the original SwRI Off-Road Test Facility DEM and the five adjusted versions of it that were used in the analysis.

Table 1 shows the parameters used to generate the five adjusted DEMs. Adjusted 1 is a baseline case, and each of Adjusted 2–5 varies one of the parameters relative to this baseline. These values were chosen to get a sense of how each parameter influences the adjustment process and metrics calculation. The $\alpha_{blending}$ parameter specifies the relative influence of the initial DEM and the adjustments in the result. When $\alpha_{blending}$ is 1.0, the result is identical to the input; when $\alpha_{blending}$ is 0.0, the result is identical to the noise.

Table 1: Parameters Used to Generate Adjusted DEMs.

DEM	Noise Variance (m^2)	$exp_{shaping}$	$\alpha_{blending}$
Adjusted 1	5.0	2.0	0.7
Adjusted 2	10.0	2.0	0.7
Adjusted 3	5.0	1.5	0.7
Adjusted 4	5.0	2.0	0.5
Adjusted 5	5.0	2.0	0.9

The blue highlighted cells in Table 1 are a visual indicator of which values are changed relative to Adjusted 1. For example, Adjusted 2 doubles the noise variance but keeps the shaping exponent and blending parameters the same as for Adjusted 1.

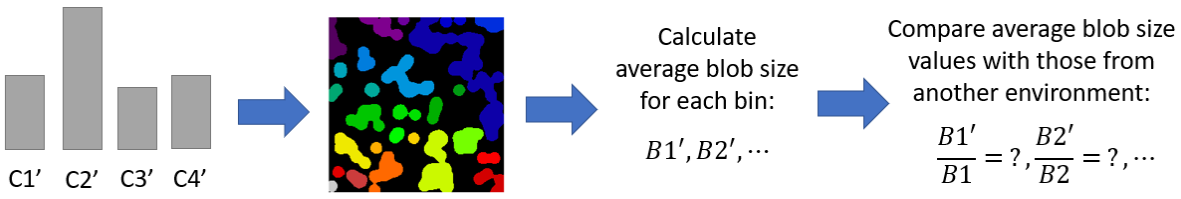

Figure 9: Flow chart of the process for the Blob Size Metric.

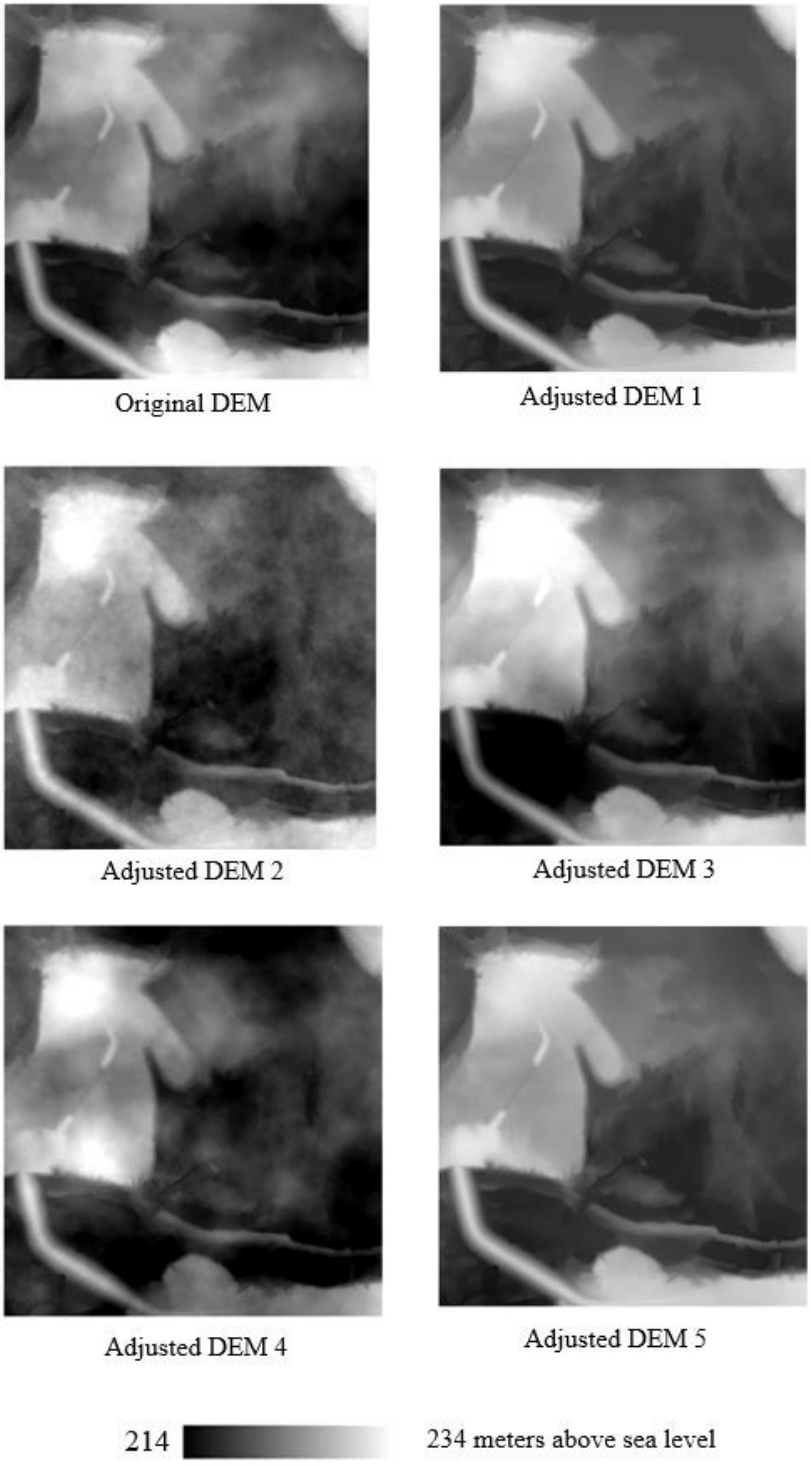

Figure 10: The original DEM and five adjusted versions used in the analysis.

For each of the six DEMs evaluated (the original and five adjusted ones), a costmap similar to Figure 7 was calculated. Cost values ranged from 0 to 255. We binned the costs into four categories: low, medium, high, and fatal. Table 2 specifies the cost bin boundaries. The same table also contains the percentages of pixels from the original DEM costmap that fall into each of the four bins.

Table 3 shows the percent changes (relative to the original) in each of the bins for the five adjusted scenes. Adjusted 1 adjusts the elevation within a variance of 5 m^2 and shapes the noise with fully red spectral content. The final blended DEM is 70% original DEM and 30% noise image. As shown in Table 3, this adjustment results in a substantial increase in the number of "medium cost" cells due to the increase in local terrain slope and roughness from the noise blending. Interestingly, Adjusted 2 (with doubled noise variance relative to Adjusted 1) does not exhibit an increase in high and fatal costs relative to Adjusted 1. This observation implies that terrain roughness, which is not greatly influenced by changes in the noise variance, has a greater overall effect on the costmap than slope.

Table 4 shows the results from the second metric, comparing the average blob sizes in each histogram bin between the original and adjusted costmaps. These results show that the DEM adjustments almost always shrink connected components. This is due to the effects of adding the noise images on the characteristics of the blobs in the adjusted images. The obvious exception in Table 4 is Adjusted 3 Bin 1, which had a dramatic increase in the average blob size relative to the original DEM. Adjusted 3 is distinguished by a shaping exponent of 1.5 (as opposed to 2.0 for all the other adjusted images). This alters the noise to be between red and pink, instead of red as for the other samples. The resulting Adjusted 3 DEM has a much higher average blob size for the low-cost regions of the terrain. When viewed visually at a high level, Adjusted 3 looks very similar to the original DEM. However, the results in Tables 3 and 4 show that the power spectral characteristics of pink noise have an outsized influence on the local shape of the terrain across the entire area, resulting in dramatically shifted costmap characteristics. The conclusion is that pink noise can be employed to strongly change the local terrain roughness.

Adjusted 4 and 5 experiment with changing the blending between the original DEM and the noise image while keeping the noise variance and pink/red noise level the same as for Adjusted 1. As expected, Adjusted 4 shows amplified effects of the noise relative to Adjusted 1, and Adjusted 5 is almost unchanged from Adjusted 1 since the weighting of the noise in the blending is only 10%.

Table 2: Costmap Histogram Bin Ranges and Original Costmap Percent of Pixels in Each Bin.

	Bin 1 (Low Cost)	**Bin 2 (Medium Cost)**	**Bin 3 (High Cost)**	**Bin 4 (Fatal Cost)**
Bin Cost Value Ranges	[0, 64)	[64, 128)	[128, 192)	[192, 255]
Percent of Pixels in Original Scene Costmap Histogram	64%	15%	13%	8%

Table 3: Percent Changes in Cost Histogram Bins of Adjusted Scenes Relative to Original Scene.

Scene	Bin 1 (Low Cost) % Change	Bin 2 (Medium Cost) % Change	Bin 3 (High Cost) % Change	Bin 4 (Fatal Cost) % Change
Adjusted 1	-35.6	+131.2	+15.0	+15.0
Adjusted 2	-29.3	+108.0	+11.2	+14.3
Adjusted 3	-100.0	-90.4	+233.1	+565.9
Adjusted 4	-78.9	+268.0	+56.2	+39.6
Adjusted 5	-2.3	+7.2	+2.0	+1.5

Table 4: Percent Changes in Average Blob Size for Each Cost Range of Adjusted Scenes Relative to Original Scene.

Scene	Bin 1 (Low Cost) % Change	Bin 2 (Medium Cost) % Change	Bin 3 (High Cost) % Change	Bin 4 (Fatal Cost) % Change
Adjusted 1	-41.0	-73.6	-23.5	-14.4
Adjusted 2	-37.2	-72.8	-18.6	-6.9
Adjusted 3	+27,357	-72.6	-59.2	-72.9
Adjusted 4	-75.9	-21.2	-66.0	-33.1
Adjusted 5	+10.5	-12.7	-0.8	-0.4

Note that the results presented here are dependent on the way costs are calculated for the costmap. Different algorithms for calculating cost will produce different quantitative values for these scene comparison metrics.

Our final metric was simply calculating the L2 (Euclidean) distance between the original and adjusted histograms. This metric provides a useful, quick, high-level indication of how similar the scenes are to one another that matches visual inspection of the scenes. However, it does not give any insight into the ways in which they differ. Table 5 shows the numerical values obtained for the five adjusted DEMs. In agreement with results from Tables 3 & 4, Table 5 shows that Adjusted 3 is the DEM that is most different from the original DEM. This is primarily due to the large increase in surface roughness due to the effects of the pink noise.

Table 5: Costmap Histogram L2 Distance from Original to Adjusted Scenes.

Scene	L2 Distance (pixels)
Adjusted 1	1.57e6
Adjusted 2	1.29e6
Adjusted 3	4.50e6
Adjusted 4	3.39e6
Adjusted 5	0.10e6

Figure 11 and Figure 12 show views of the 3D scenes, comparing the SwRI Off-Road Test Facility generated from the original GIS data with scenes from the adjusted DEMs and foliage mask images. Figure 11 shows how surface roughness can be altered using the

DEM adjustment techniques. Figure 12 demonstrates how adjusting the foliage mask image can change the layout and density of the trees and grasses.

Figure 11: Top: a view of a berm in the SwRI Off-Road Test Facility with the original DEM. Bottom: a similar view with the Adjusted 3 version of the DEM, showing an increased surface roughness.

4. DISCUSSION AND CONCLUSIONS

This research contributed a novel software pipeline to generate realistic 3D simulation scenes from real-world GIS data, tools to intelligently adjust relevant geo-environmental variables, and novel metrics to compare scenes to one another based on characteristics relevant to off-road AV systems.

This research advances the state of the art in simulation generation for off-road AV development. The results show that, in general, the adjustment tools serve to increase navigational costs for the vehicle and decrease the connected component sizes of regions of similar cost.

Figure 12: Top: a view in the 3D scene generated from the Adjusted 3 DEM and the original tree and tall grass mask images. Bottom: view from the same vantage point with the same DEM but adjusted tree and tall grass mask images causing different foliage arrangements and densities.

The tools provide a quantitative basis for an intuitive understanding of how varying each of the parameters in the adjustment process changes the end results, but more research is needed in this area. Further research would enable the creation of an end-user system with intuitive tools that are abstracted from these parameters, such as “make the foliage denser,” “add more steep drop-offs to the terrain,” or “add more water obstacles.” Such a system would accelerate autonomy and robotics research by bypassing expensive and time-consuming scene generation and modification by simulation experts and artists and giving direct control to the researchers.

Additional valuable adjustments to the scene would include the following:

- Decrease navigation cost by flattening and smoothing the terrain.
- Increase the average connected component size of regions of similar cost, i.e., generate larger forests, bigger grassy fields, or wider dirt paths.
- Make larger-scale modifications to the terrain, such as adding or subtracting complex features like berms, ravines, hills, and water features. Effectively parameterizing these adjustments would enable the creation of completely new scenes that nonetheless share important characteristics with the geo-specific scene of interest.
- Add/subtract additional types of foliage and vary the size of the foliage.

The portfolio of metrics for comparing scenes to include measures of larger-scale and more complex terrain features such as those described above can also be expanded.

5. REFERENCES

[1] J. Brabbs, S. Lohrer, P. Kwashnak and M. Brudnak, "M&S as the Key Enabler for Autonomy Development, Acquisition and Testing," in *Proc. of the Ground Vehicle Systems Engineering and Technology Symposium (GVSETS)*, 2019.

[2] J. Boone, C. Goodin, L. Dabbiru, C. Hudson, L. Cagle and D. Carruth, "Training Artificial Intelligence Algorithms with Automatically Labelled UAV Data from Physics-Based Simulation Software," *Applied Sciences,* vol. 13, no. 131, 2023.

[3] "Unreal Engine," Epic Games, [Online]. Available: https://www.unrealengine.com/en-US/.

[4] "San Antonio River Authority," [Online]. Available: https://www.sariverauthority.org/.

[5] "OpenStreetMap," [Online]. Available: https://www.openstreetmap.org/.

[6] "USGS EarthExplorer," [Online]. Available: https://earthexplorer.usgs.gov/.

[7] "QGIS: A Free and Open Source Geographic Information System," [Online]. Available: https://qgis.org/en/site/.

[8] "GIMP: GNU Image Manipulation Program," [Online]. Available: https://www.gimp.org/.

[9] Environmental Systems Research Institute, Inc, "ESRI Shapefile Technical Description," 1998.

[10] "TerraForm PRO," [Online]. Available: https://www.terraformpro.com/.

[11] A. Patel, "Making Maps with Noise Functions," May 2020. [Online]. Available: https://www.redblobgames.com/maps/terrain-from-noise/.

[12] P. Bourke, "Generating Noise with Different Power Spectra Laws," October 1998. [Online]. Available: http://paulbourke.net/fractals/noise/.

[13] H. Tulleken, "Poisson Disk Sampling," 3 May 2009. [Online]. Available: http://devmag.org.za/2009/05/03/poisson-disk-sampling/.