Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
May 2, 2024
A Regional Conference of the Society for Industrial and Systems Engineering

Modeling Terrain Effects on Non-Strategic Nuclear Weapons

Michael O'Connor, Zachary Dugger, and Alexander Withenbury

Department of Mathematical Sciences,
United States Military Academy,
West Point, NY 10996

Corresponding author's Email: michael.j.oconnor316.mil@army.mil

Author Note: Michael O'Connor is a First Class cadet (senior) at the United States Military Academy at West Point and is pursuing a Bachelor of Science degree in Mathematical Sciences. Following graduation, CDT O'Connor will commission as a Second Lieutenant in the Signal Corps. MAJ Alexander Withenbury and MAJ Zachary Dugger are Assistant Professors in the United States Military Academy's Department of Mathematical Sciences. We would like to thank LTC Nickolas Duncan and the Defense Threat Reduction Agency for their support of this research.

Abstract: Building upon previous work, this project improved an existing simulation program for the tactical employment of non-strategic nuclear weapons by improving its predictive damage equations and incorporating shielding effects of terrain. Using data from the United States Geological Survey, this research created a sub-routine within the model to determine if there is a line of sight between detonation location and a potential target, adjusting the predicted damage accordingly. We found that this addition could have a significant impact on casualty estimates. However, the increased runtime had a detrimental effect on the usability of the estimate. The improved model will provide battlefield commanders with more accurate casualty estimates during wargames and will increase mission readiness for a nuclear battlefield. Further research should focus on increasing the efficiency of the simulation.

Keywords: Non-Strategic Nuclear Weapons, Wargaming, Simulations, Viewshed

1. Introduction & Background

In an increasingly unstable global environment, the use of nuclear weapons on the battlefield is an ever-present threat. Recent events, such as the Russian invasion of Ukraine and heightened conflict in the Middle East, have raised the prospect that low-yield nuclear weapons may be used on the battlefield (Woolf, 2022). Though there are significant policy questions concerning the use of such weapons, this research focused narrowly on the effects of non-strategic nuclear weapons (NSNWs) on military units. Definitions vary regarding what constitutes an NSNW; however, it is considered to be a weapon with a yield of 30 kilotons or less for the purposes of this project (Woolf, 2022). Research indicates that the employment of nuclear weapons with limited yield carries lower risk of broader escalation to a full-scale nuclear conflict and would be used in lieu of more destructive weapons, though there is some disagreement among scholars, as well as a lack of empirical data (Reddie and Goldblum, 2022). Further questions about the potential employment of nuclear weapons are interesting and important, but they are beyond the scope of this research.

This research builds on previous work that created a wargaming program using a Monte Carlo simulation to allow commanders to accurately assess the effect of an NSNW on their units (Guetzke, 2023). This previous work incorporated extensive details of Army doctrine and open-source materials to simulate the effects of a nuclear weapon based on detonation location and Soldier positioning on the battlefield. The simulation outputs easily interpretable data depicting the expected casualty rate (modeled from 1,000 Monte Carlo iterations), distinguishing between Soldiers which are killed, injured, and uninjured, and vehicles which are destroyed, damaged, and undamaged.

To improve the accuracy of the existing model, this project sought to refine the implementation of radiation damage and incorporate terrain effects on the simulation. Nuclear explosions inflict damage through a pressure wave (or "blast"), thermal radiation, and gamma and neutron radiation (Urone and Hinrichs, n.d.). Though each of these effects was accounted for in the model created by Guetzke et al., this project aimed to improve upon the radiation and thermal damage calculations by focusing specifically on the impact of terrain on radiation and thermal damage. The interaction between the blast wave and terrain is very complex and thus remained outside the scope of this research. Overall, this project aimed to improve the simulation's ability to accurately represent the thermal and radiation effects of the weapon, both of which would be reduced by the presence of a terrain feature between the weapon and the target.

2. Methodology

2.1 Existing Model

The model created by Guetzke et al. features a graphical interface that allows the user to provide several inputs. The program then uses a Monte Carlo simulation to return results and visualizations to the user. Through the interface, which is depicted in Figure 1, the user controls several aspects of the nuclear weapon employment. First, the user enters the yield of the weapon, up to the 30-kiloton limit imposed on the model, and then selects one of several height of burst options: ground burst, fallout-free (which effectively eliminates the long-term radiation effects on the environment), optimal for 50 pounds per square inch (psi) of peak overpressure, and optimal for 10 psi of peak overpressure (Northrop, 1996). The values of 50 psi and 10 psi were chosen because these values correspond to roughly 50% lethality in humans and the threshold for injury if standing, respectively (Glasstone, 1977). Next, the user selects the coordinates of the NSNW's detonation. Lastly, the user selects the delivery mechanism for the NSNW—either deployed from an intercontinental ballistic missile, bomber, or submarine. Each of these selections corresponds to an estimated error value reflecting the expected precision of the delivery method. For a more detailed discussion of these values, see Guetzke et al (2023).

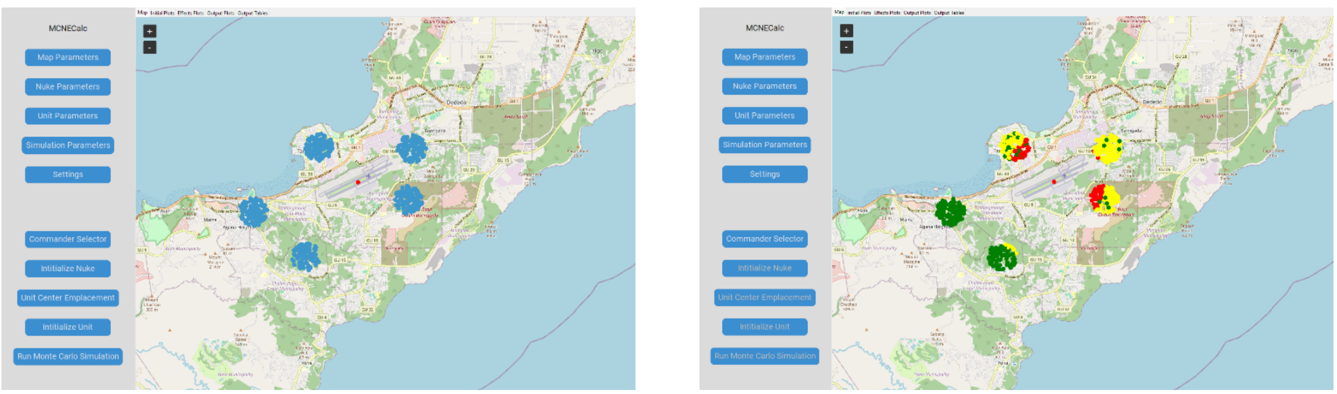

(a) The user's inputs before executing the simulation.

(b) The results displayed to the user after executing the simulation.

Figure 1. An example of the user interface of the simulation created by Guetzke et al. and built upon during this research.

2.2 Modeling Terrain

In incorporating terrain effects into the model, this project only considered whether the terrain blocks the ray between the NSNW and the target—if the terrain does block this ray, we consider the target to be shielded from the thermal and radiation effects of the NSNW, and if the terrain does not block the ray, then we consider there to be no effect from terrain. The peak overpressure, meanwhile, is unaffected by terrain in our model. Notably, these are oversimplifications; in particular, radiation can refract over obstacles in a phenomenon known as skyshine, and the wave of peak overpressure will have a complicated relationship with terrain, as its intensity at a given point may be changed by the slope of the terrain there and in the surrounding area (McDermott, 2020; Northrop, 1996). However, the computation of these effects is significantly more difficult and is beyond the scope of this research. The focus of this project was limited to examining if a simplified terrain model could substantially alter the simulation's predictions to more closely reflect the physical reality.

Publicly available data was used to create the terrain model for this simulation. The United States Geological Survey (USGS) offers vast amounts of topographic data of varying resolutions from areas around the United States (3D Elevation Program, n.d.). The scope of this project was limited to the island of Guam, in order to validate the incorporation of terrain on a relatively small area. This location was selected for its strategic importance in the Pacific region, the presence of U.S. military installations on the island, and the varying geography for testing terrain effects. This allowed us to use only one raster (a matrix of geographic data) from the USGS for our implementation and testing.

A difficulty associated with considering terrain in the model was to do so while still allowing the model to run in a tactically feasible amount of time (defined as fifteen minutes for this research). If the simulation required more time than this, it would affect its usability and decrease its value for wargaming efforts. At the outset of this research, the simulation could already take several minutes to run on a typical computer available to a commander, depending on the parameters selected by the user. Since a large elevation raster can be time-consuming to access, and many calculations are required for a single line-of-sight analysis, runtime would be a persistent challenge throughout the development of the model.

3. Model

3.1 Adjusted Damage Calculations

This research adjusted the damage calculations used in the simulation for the radiation and thermal effects to improve the accuracy of the simulation. Each of these were adjusted using estimates and equations from Northrop (1996). To validate the changes, the results were compared to those of the Mission Impacts of Nuclear Events Software (MINES), the program of record for nuclear weapons simulations created by the Defense Threat Reduction Agency (DTRA). The improved equations yielded similar results to MINES, suggesting their validity and accuracy. Exact results are not reported due to the classification level of MINES.

The program of record does not currently implement a terrain model in its calculations. Therefore, only raw damage outputs were assessed against MINES, and the results from terrain implementation were not tested for accuracy against MINES.

3.2 Line of Sight Calculations

The central task in incorporating terrain into the model was determining whether there was a line of sight between two given points. Detecting the existence of such a line of sight would allow the model to determine if a terrain feature was shielding the target from the thermal and radiation effects of the NSNW. The model would therefore need to perform this calculation thousands of times, quickly enough to be useful for wargaming. Thus, the algorithm needed to be simple and fast.

The process of determining if there is an uninterrupted line of sight between two points is well-studied, with existing solutions such as ray tracing and viewshed analysis, depending on the application. Ray tracing is an algorithm that follows the reflections of a theoretical light ray as it travels and reflects to a light source. The technique is often used in computer graphics to generate realistic depictions of shadows and reflections (Nvidia, n.d.). Our implementation of terrain, then, amounts to a simpler version of this algorithm. Rather than tracing a given ray through a series of reflections off obstacles until it terminates, we are only interested in the first intersection (if it occurs at all before the ray terminates at the other endpoint). This means that, although ray tracing is a powerful algorithm, it does not provide a ready answer for approaching the difficult step of determining if there is a line of sight—namely, calculating precisely where the ray intersects with an irregular surface.

A more directly applicable field of study is viewshed analysis. Viewshed analyst tools, such as that included in the software ArcGIS Pro, allow the user to determine which areas of a map are visible to an observer at a given height above a certain location (Esri, n.d.). Due to the inherent irregularity of real terrain surfaces, and subsequent infeasibility of using analytic techniques, viewshed analysis conventionally employs a numerical approach. The algorithm will start at the observer point and take incremental steps toward the target point. At each point, the program evaluates the height of the ray and the height of the terrain at that point. If at any point the elevation of the terrain exceeds the elevation of the ray, the algorithm terminates and returns that there is no line of sight. If the algorithm reaches the target point without this occurring, then the algorithm terminates and returns that there is a line of sight (Achilleos and Tsouchlaraki, 2004).

This algorithm results in thousands of calls to the data structure that stores the elevation data. Due to the inherent slowness in accessing a large raster, direct implementation of this algorithm can lead to an unacceptably long runtime, depending on what step size is used. In order to shorten the runtime while preserving accuracy, the simulation uses a shortcut that is effective for certain NSNW locations and heights of burst. First, each line of sight begins tracing from the target and terminates at the bomb. It then uses a specified upper bound on the elevation of terrain in the area under consideration—that is, the maximum altitude at which the line of sight could be cut off by an obstacle. If the ray reaches this altitude and has not encountered an obstacle, the run terminates, returning that there is a line of sight. This allows it to cut out many superfluous calls to the elevation raster for higher heights of burst, shortening the runtime in those cases. However, the runtime of this process is still a significant problem and is an important area for further study.

4. Results

Testing confirmed that the shielding effect of terrain has a significant impact on the model in many cases. In the example pictured in Figure 2, an Infantry Company was predicted to be at 56.06% combat strength after the NSNW detonated when using the original model. In contrast, when terrain was incorporated, it was predicted to be at 82.34% strength. In each test, the NSNW and the Infantry Company were centered at the same locations on the map. Note that each Soldier in subplot (a) has an elevation of 0 meters, while the Soldiers in subplot (b) are placed at an elevation corresponding to the terrain beneath them. Additionally, the NSNW in subplot (b) is placed almost 100 meters higher than the NSNW in (a) because its height of burst is measured from the height of the terrain, rather than sea level. In both cases, a large group of green points suggests that they are outside the injuring radius of the NSNW. Outlying green points amidst yellow points are likely vehicles, which have higher damage thresholds.

Over the 1,000 iterations of the Monte Carlo simulation, the NSNW and the Soldiers within the Infantry Company were randomized to new locations—the NSNW according to a multivariate normal distribution, and the company uniformly distributed within a disk centered at the selected coordinates. Therefore, each test was performed on a slightly different set of unit and bomb positions, but the average of the Monte Carlo simulation is expected to converge to the true value over 1,000 iterations. Additionally, the effects of terrain could be much larger or smaller in different circumstances—in flat terrain, the effect will be negligible, while terrain could completely block the thermal and radiation effects of the NSNW if the unit is located on the opposite side of a ridge.

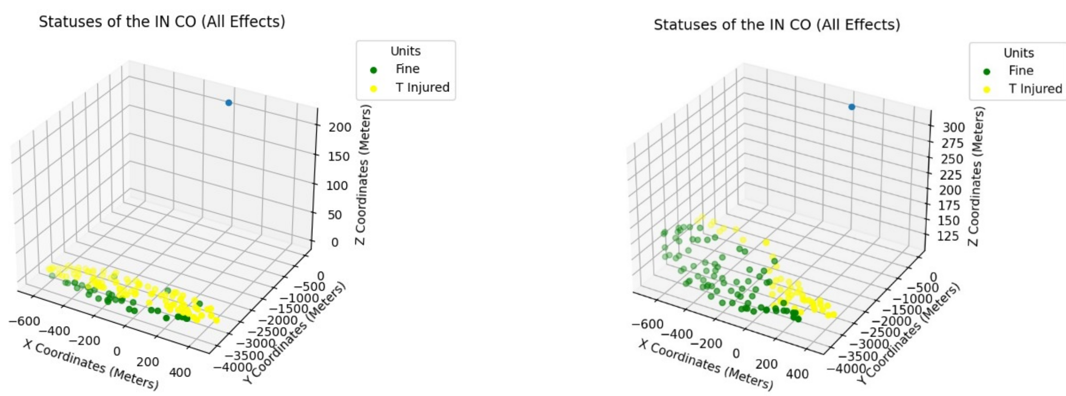

(a) Without accounting for terrain. (b) Accounting for terrain.

Figure 2. Plots of the effects of a 30 kiloton NSNW. The blue point indicates the point of detonation of the NSNW, yellow represents those injured by thermal effects, and green indicates uninjured Soldiers. Only thermal effects caused injuries.

4.1 Performance Comparisons

Since both the line of sight and damage calculations are performed separately for each individual Soldier and vehicle, the runtime of the model increases linearly with the number of Soldiers and vehicles modeled. This is true both when accounting for terrain and when terrain is excluded from the model, as seen in Figure 3. The simulations that generated this data were performed on a Dell Latitude 5400 with an Intel® Core™ i7-8665U CPU processor with 32 GB of dual-channel DDR4 RAM. The high R^2 value of both lines indicates that the number of points being modeled almost entirely explains the runtime of the model. The random variation caused by available CPU resources, for instance, is not significant in comparison. Points were sampled from six units of a range of sizes that the simulation can model. The largest unit, an Infantry Brigade Combat Team, is modeled with 4421 Soldiers and vehicles, while a Field Artillery Company within an Armor Brigade Combat Team is the smallest, with 73 Soldiers and vehicles.

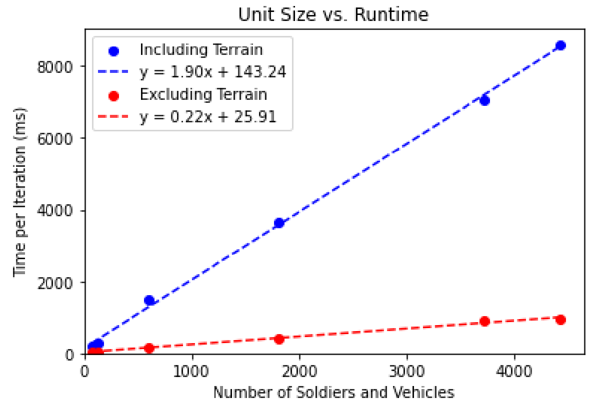

Figure 3. Plot depicting the runtime of the model for units of varying sizes, both including terrain and not including terrain. For the model including terrain, $R^2 = 0.998$, and $R^2 = 0.992$ for the model without terrain.

The slope of the line fitting the model with terrain is approximately 8.62 times the slope of the other line, showing that modeling a given unit will take almost nine times more time with terrain than without terrain. The respective slopes tell us that each additional Soldier or vehicle increases the runtime by 1.90 milliseconds per iteration when terrain is modeled, and 0.22 milliseconds per iteration when terrain is not modeled. These precise numbers can be expected to vary widely depending on the processor running the code. However, the linear relationship between unit size and computation time is expected to remain, as is the longer runtime associated with the terrain model.

This extra computational cost is significant for larger units and numbers of iterations, with an Infantry Brigade Combat Team taking two hours and twenty-three minutes to run 1000 Monte Carlo Iterations. This suggests that commanders may wish to use the model containing terrain when modeling smaller units or when computation time is not a limitation. For larger units, however, commanders may instead use the model without terrain if they do not believe there are sufficient terrain features to impact the results. Alternatively, future research could identify a smaller, optimized number of Monte Carlo iterations.

4.2 Monte Carlo Convergence

Over a large number of iterations, a Monte Carlo simulation is expected to converge to the true population parameter. However, in cases where each iteration is computationally expensive, we are interested in calculating how many iterations yields a sufficiently precise answer. Oberle (2015) gives the following formula for estimating the necessary number of iterations n in the case of a binomial distribution:

$$n = \frac{Z_{\alpha}^{2}\hat{p}^{2}\left(1-\hat{p}\right)}{\Delta^2} \quad `(1)$$

where Z is the Z-score, α is the significance level, $\hat{p}$ is the observed proportion of successes (used as an estimator for the population parameter), and Δ is one-half of the desired confidence interval. Calculating n with $\alpha = 0.05$ and $\Delta = 0.05$ (so the 95% confidence interval of the percent strength of the unit will be approximately 10% wide), and assuming the worst-case scenario of $\hat{p} = 0.5$, yields $n = 385$ iterations. This suggests that, by making 385 the default number of iterations rather than 1,000, the runtime of the simulation can be significantly reduced without sacrificing accuracy. Further research should empirically examine the convergence of this model, confirming this estimate or finding an alternative optimal number of iterations.

5. Conclusions & Discussion

The implementation of terrain in this model shows that it can have a substantial impact on the predicted effects of an NSNW. This could have implications for commanders of any echelon who wish to predict the possible effects of an NSNW. Uneven terrain could be highly advantageous for the survivability of forces, and accounting for terrain effects will yield more accurate predictions of the impact on friendly forces. As commanders face the possibility that these weapons will be deployed

ISBN: 97819384962-4-0

on the battlefield, a greater understanding of how terrain can be used to protect friendly forces is critical. Therefore, the improvements made in this research provide valuable information as compared to models that do not incorporate the shielding effects of terrain.

The most significant drawback of these changes is the increased computational cost of the model. For commanders of smaller units such as companies, this difference is less pronounced, since the simulation completes in a feasible amount of time regardless of whether terrain is included. Higher echelon commanders, meanwhile, must allow significantly more time for the simulation if they wish to account for terrain. However, the option to include terrain effects is still valuable for commanders of large units and gives them the ability to improve the simulation's accuracy. Moreover, commanders may choose a smaller number of iterations in order to sacrifice some certainty of the casualty estimate to reduce runtime.

Though there are many remaining ways to increase the fidelity of this model (such as including dynamic overpressure or long-term radiation), the current computational cost of the simulation may render such improvements impossible until the algorithm's efficiency is improved. Therefore, future work should focus on the difficulties associated with the runtime of incorporating terrain into a large-scale simulation, including the time required to access the elevation data. This could include using a more efficient method of data storage or the use of dynamic programming to avoid repetitive calls to the raster. Additionally, researchers could use a more efficient method of viewshed analysis, replacing the naïve, brute-force approach of taking incremental steps along a ray.

6. References

3D Elevation Program. *What is 3DEP?* U.S. Geological Survey. Retrieved from: https://www.usgs.gov/3d-elevation-program/what-3dep.

Achilleos, G. & Tsouchlaraki, A. (2004). *Visibility and viewshed algorithms in an information system for environmental management*. Management Information Systems. Retrieved from: https://www.witpress.com/Secure/elibrary/papers/MIS04/MIS04012FU.pdf.

Esri. *Viewshed (Spatial Analyst)*. Esri. https://pro.arcgis.com/en/pro-app/3.1/tool-reference/spatial-analyst/viewshed.htm

Glasstone, S. & Dolan, P. J. (1977). *The Effects of Nuclear Weapons*. United States Department of Defense.

Guetzke, T., Withenbury, A., Dugger, Z. & Kendall, T. (2023). Modelling Nuclear Weapon Effects in Wargaming Using Monte Carlo Simulations. *Proceedings of the Annual General Donald R. Keith Memorial Conference*.

McDermott, P. N. (2020). Photon skyshine from medical linear accelerators. *Journal of Applied Clinical Medical Physics, 21*(3), 108-114. doi:https://doi.org/10.1002/acm2.12833.

Northrop, J. A. (1996). *Handbook of Nuclear Weapon Effects*.

Nvidia. *Ray Tracing*. Nvidia Developer. https://developer.nvidia.com/discover/ray-tracing.

Oberle, W. (July 2015). Monte Carlo Simulations: Number of Iterations and Accuracy. *US Army Research Laboratory*.

Reddie, A. W., & Goldblum, B. L. (2023). Evidence of the unthinkable: Experimental wargaming at the nuclear threshold. Journal of Peace Research, 60(5), 760-776. Retrieved from: https://doi-org.usmalibrary.idm.oclc.org/10.1177/00223433221094734.

Urone, P. R. & Hinrichs, R. (n.d.). *College Physics*. OpenStax. Retrieved from: https://openstax.org/books/college-physics/pages/preface.

Woolf, A. F. (2022, March 7). *Nonstrategic Nuclear Weapons.* Congressional Research Service.

ISBN: 97819384962-4-0