

A General Purpose Geospatial Encounter Prediction Model for Border Security

Allen S. Harvey Jr., Damian Kolbay, Jesse Coleman, & Jessica McNutt

Engility Corporation

Lorton, VA

**allen.harvey@engilitycorp.com, damian.kolbay@engilitycorp.com,
jesse.coleman@engilitycorp.com, jessica.mcnutt@engilitycorp.com**

ABSTRACT

The Department of Homeland Security (DHS) is interested in understanding capabilities to encounter people crossing the borders between official ports of entry, and in determining the likelihood and probable areas where law enforcement officials can intercept a suspect, either on or off road. Engility Corporation, under contract with DHS, has developed a general purpose geospatial encounter model as a plugin for QGIS. This model takes into account law enforcement and adversary vehicle types, road networks, natural boundaries, elevation, and land type to quickly and efficiently generate triangular meshes to compute isochrones accurately over entire regions. These isochrones support a physics constrained, quantitative assessment of potential encounter capabilities. Novel modeling approaches were introduced to improve computational efficiency of modeling multiple law enforcement and adversary entities in a single scenario. The software can also aid in planning future force structure laydowns to detect illegal border crossings and off road searches.

ABOUT THE AUTHORS

Allen Harvey is a Computational Scientist Associate Technical Fellow at Engility Corporation. He obtained his Bachelor's degrees in Mathematics, Computer Science, and Physics at SUNY Brockport and his Masters in Computational Science at George Mason University. He is currently a PhD Candidate in Computational Physics at George Mason University and a certified Project Management Professional by the Project Management Institute. He leads Engility Corporation's modeling and simulation efforts supporting the Domestic Nuclear Detection Office.

Damian Kolbay is a Modeling and Simulations Engineer at Universal Solutions International, Inc. (supporting Engility Corporation) with an extensive background in geospatial analysis and modeling. He obtained his Bachelor's degree in Geography at the University of Tennessee and his Master's degree in Geographic and Cartographic Sciences at George Mason University.

Jesse Coleman is a Requirements Analyst and Systems Engineer at Engility Corporation. He obtained his Bachelor's and Master's degrees in Civil, Environmental, and Infrastructure Engineering at George Mason University. He is an engineer in training (EIT), and his expertise includes land development, geographic information systems, and 3-D modeling,

Jessica McNutt is a Modeling and Simulations Engineer at Engility Corporation. She obtained her Bachelor's degree in Physics and Mathematics from the University of North Carolina at Asheville and her Master's degree in Nuclear Engineering from the Pennsylvania State University.

A General Purpose Geospatial Encounter Prediction Model for Border Security

Allen S. Harvey Jr., Damian Kolbay, Jesse Coleman, & Jessica McNutt

Engility Corporation

Lorton, VA

allen.harvey@engilitycorp.com, damian.kolbay@engilitycorp.com,
jesse.coleman@engilitycorp.com, jessica.mcnutt@engilitycorp.com

BACKGROUND

Customs and Border Protection's (CBP) Border Patrol office is responsible for responding to any unauthorized incursion across the borders. Sensor systems, such as cameras (visible and infrared), acoustic sensors, and motion sensors, have proven effective at discovering incursions (U.S. Customs and Border Protection, 2011). These systems may be deployed on different platforms including land, towers, blimps, vehicles, planes, and helicopters. Once law enforcement discovers and encounters a suspect, additional means may be used to determine the presence of any illicit materials that should not be entering the country. Thus, the probability of detecting any of these materials from a border incursion can be characterized as a probability of discovery and encounter problem (Edward, 2013 & Willis, 2010).

To understand the abilities to encounter entities crossing the borders illegally, DHS required the creation of a model that would allow an analyst to assess changes in capabilities given varying definitions of the threat, sensors, and CONOPs for both law enforcement and adversaries along the entire Northern and Southern Border regions simultaneously and holistically. Unfortunately, access to specific CONOP data used by law enforcement would not be available during the timeline for this project.

APPROACH

To address the need to assess the entire border regions without access to law enforcement CONOPs, Engility Corporation approached the problem with a physics constrained, geospatial perspective. This approach looks at the best case scenario for law enforcement to respond to an alert of a border incursion, given capabilities of their response assets (e.g., cars, all-terrain vehicles, horses, etc.) to travel on-road and off-road given the land type, terrain elevation, and natural barriers. At the same time, an adversary would be subject to the same conditions as it crosses anywhere along the border and heads inland. This approach compares the areas that an adversary could reach and the overlap in area that law enforcement could reach to assess a physics-constrained probability of encounter.

THE GEOSPATIAL ENCOUNTER MODEL (GEM)

The Geospatial Encounter Model (GEM) draws upon a wide variety of open source and proprietary data and uses several different computational algorithms to determine the areas within the U.S. Border region law enforcement and adversaries can reach in a given period of time. This model has been constructed as a plug-in for the free and open source geospatial information systems software, QGIS. QGIS provides a rich environment to conduct a wide range of modeling, simulation, and analysis relying on industry standard and community supported technologies and techniques. By being free and open source, there are no licensing restrictions or costs, allowing Engility Corporation to rapidly construct and deliver a tool for DHS to use. The software provides an open applications programming interface via Python allowing Engility Corporation to automate workflows, provide a Graphical User Interface (GUI), and develop the model. The required data and the tool's methodology are described below.

Data Sources

A large quantity of data was gathered during GEM's development phase. Broadly speaking, the data falls into two distinct categories. The bulk of the data – terrain, land-use, road networks – is geospatial data. This data is essential for the analysis of pathways, both on and off-road, as performed by the tool. The rest of the data is blue force asset data and includes information regarding the location of sensor mounted towers, unattended ground sensors,

checkpoints, ports of entry, etc. All of the data essential for operation of the tool is included with the software package.

Terrain

Terrain data was collected from an online database, called National Map Viewer, maintained by the U.S. Geological Survey's National Geospatial Program (United States Geological Survey, 2014). The database provides information regarding elevation, orthoimagery, hydrography, geographic names, boundaries, and U.S. topographical and historical maps. National Map Viewer provides elevation data in many different resolutions and file formats and is compatible with most geospatial analysis software, including QGIS. The data selected for this analysis was one arcsecond resolution for .IMG files. Files were downloaded to cover the entire Northern Border with Canada and the Southern Border with Mexico.

Land Type

Land use data was taken from the National Park Service and the National Land Cover Database (NLCD) hosted by the Multi-Resolution Land Characteristics Consortium (United States Geological Survey & United States Department of the Interior, 2014 & National Park Service, 2014). Land use data was used solely for the analysis of terrain traversal (discussed later). The Land Type data from NLCD classifies the CONUS into 16 categories (4 additional categories are Alaska only) at 30 meter resolution. The land use map of the CONUS and the (non-Alaska) categories are shown in Figure 1.

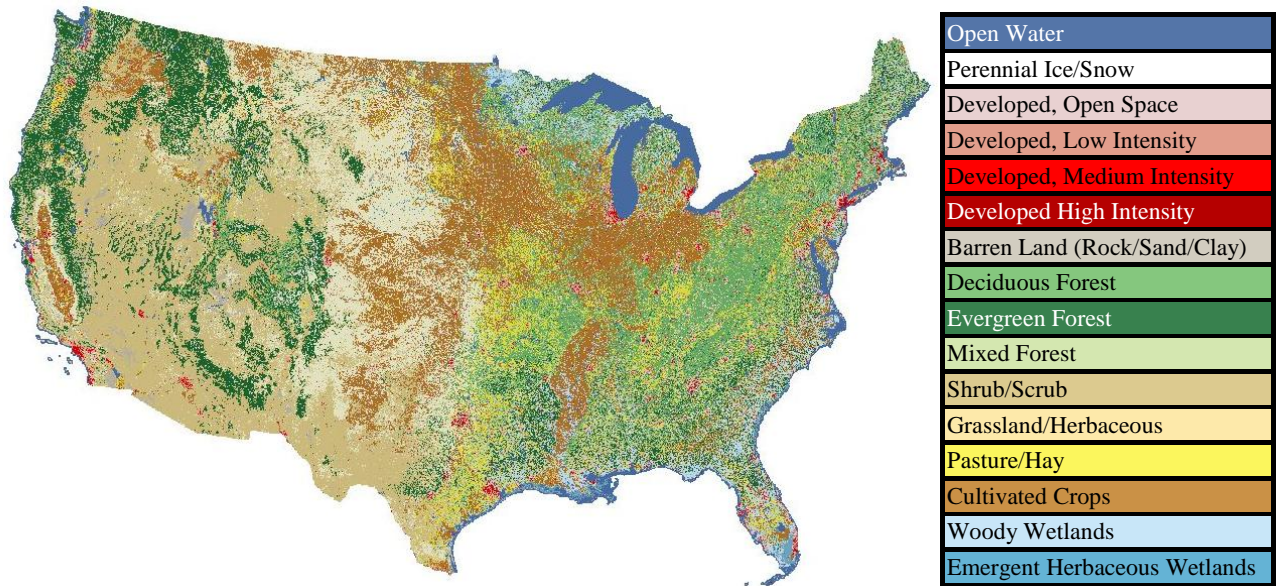


Figure 1. Land Use Map of the CONUS

Urban Area Data

Population data, available through the U.S. Census, was used to determine the location of urban areas throughout the Northern and Southern Border regions (U.S. Department of Commerce, 2014). GEM allows the analyst to treat urban areas as win zones, meaning that blue forces would be unable to encounter red forces within the urban area zones, thereby, reducing the overall probability of encounter. For some scenarios, it may be appropriate to restrict law enforcement from being effective or able to operate within an urban environment and vice versa. Additionally, these urban areas can be used for visualization purposes, to aid the analyst in determining scenario times, sensor delay times, etc. Alternatively, the tool allows the user to specify any vector layer containing polygons to be used as a way to reduce blue force effectiveness.

Road Network Data

Road and highway data used for this analysis was taken from a free, crowd-sourced wiki called OpenStreetMap (OpenStreetMap, 2014). OpenStreetMap is maintained by a community of online mappers and includes everything from roads, railroads, highways and trails to restaurants, houses and more. While the “wiki” nature of this data

source can potentially lead to errors, the number of contributors and the high degree of transparency in all map additions leads to a highly reliable map, especially when used on a large scale. Although much of the data is crowd sourced, the base of the data is provided by the US Census Topologically Integrated Geographic Encoding and Referencing (TIGER) data. This data can be downloaded directly from the OpenStreetMap website (<http://www.openstreetmap.org>) or through the QGIS OpenLayers plugin available within the QGIS software. GEM only uses a subset of the available OpenStreetMap data, including road networks, waterways, and barriers.

Pathway Network Analysis and Mesh Generation

To construct an off-road network for adversaries and law enforcement to travel on, the border region was meshed using the two-dimensional mesh generator called Triangle (Shewchuk, 2005). Triangle is a free and open source Delaunay triangulation mesh generator developed by Carnegie Mellon University as part of an initiative to improve large-scale earthquake simulations. This software expands on the OpenStreetMap input networks by adding points to a region and effectively creating a complete network of on-road and off-road pathways through the creation of new edges obtained from meshing. Mesh generation using the Yuma sector road network is shown in Figure 2. By using a triangle mesh generator, GEM preserves the road network geometry and avoids the introduction of square edge anomalies common when using alternative Cartesian mesh generators. A mesh for each CBP sector was created and is distributed with the model. The meshes provided with this model maximized resolution subject to a constraint to limit model runtime to the order of minutes.

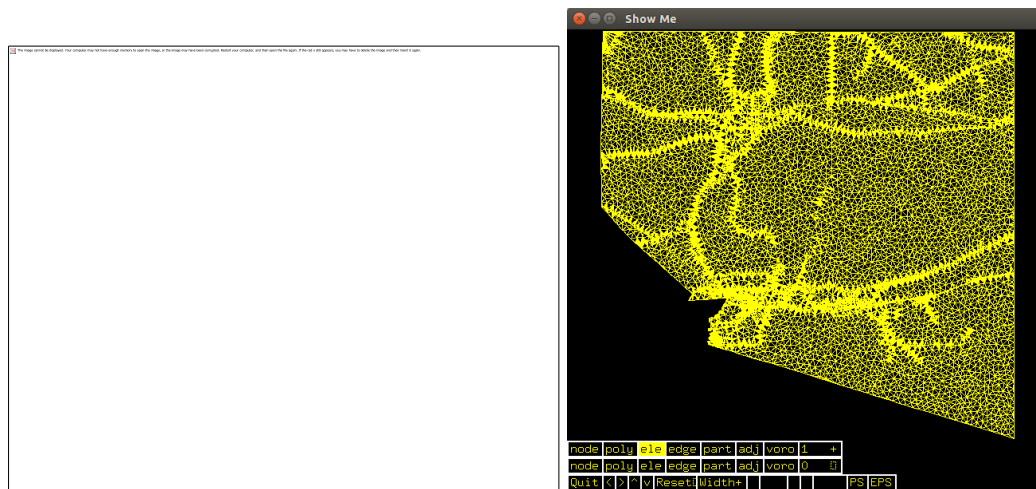


Figure 2: Yuma sector road network and urban area before (left) and after (right) meshing

Due to the large size of meshes containing on-road, off-road, and barrier features, each sector was generated as a standalone mesh. The standalone mesh approach was required due to limited computer memory available for processing travel time calculations (standard workstations used for this project had 8 gigabytes of system memory available). This approach has both positive and negative impacts on the modeling environment. During calculation of the travel graphs (described later), the calculation time is significantly faster due to the reduced size of the analysis area. A negative result of the sector approach is the inability of a force to travel between sectors. If a red or blue force is traveling in one sector, GEM currently does not allow entities to cross into an adjacent sector. Depending on CBP CONOPs, this approach of not implementing travel between sectors may be appropriate for blue forces, but is certainly not advantageous for modeling red forces. In an effort to bridge the gap within GEM, a force that is seeded within three kilometers of a sector border will be seeded in both sectors.

Travel Cost Graph

Once the mesh network is output from Triangle, costs are assigned to each edge as a measure of the speed in which an entity might traverse that unit of distance. These costs are generated using a variety of factors, such as the presence of roads or highways, barriers, land classification, slope, and vehicle or conveyance type. For example, a flat stretch of road would receive a higher speed than a steep incline of mountain with no road. The level of detail for this stage of the analysis is dependent on the resolution of the underlying mesh used to define the region and the resolution in which edges are sampled against land type and terrain elevation data to determine an average speed over the edge (see Figure 3). GEM provides a fixed set of meshes, but the analyst can choose between several

different mesh sampling resolutions. Higher resolution edge sampling provides more accurate travel costs at the expense of increased processing time.

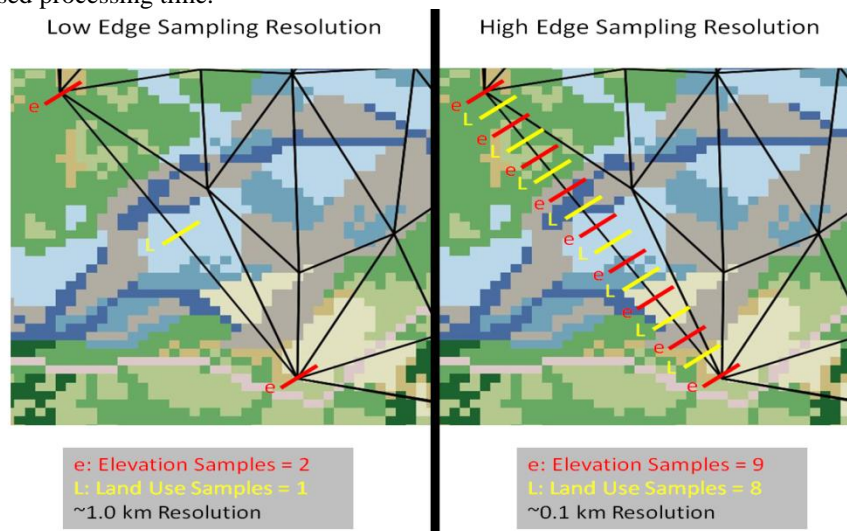


Figure 3. Edge Sampling Resolution Comparison

The creation of the travel graph is the most computationally expensive part of the model. This is due to the need to query millions of points in the raster land type and elevation files. To improve model performance, each time a user runs a scenario using a new resolution and vehicle type combination, the resulting cost graph is saved to the disk and cached for future reuse. This way, if the same scenario is run again in the future, this step will take no time at all. Should the user change the underlying meshes, elevation, vehicle speed tables, or land type data objects, then this cached data must be flushed and regenerated.

Shortest Path Calculations

The distances an entity can travel within the border regions is achieved by using a priority queue optimized Dijkstra's shortest path algorithm with $O(m \log n)$ complexity (with 'm' being the number of edges in a graph and 'n' the number of nodes). Normally, this algorithm is used to evaluate the single shortest path from points A to B. However, in computing this path, the algorithm evaluates the distance to every point. This means, by finding the shortest path to any arbitrary node in the graph, the shortest path to all nodes is found, too. Although Dijkstra's algorithm is typically associated with finding a distance, if using weights, it can find the lowest cost pathway to reach a node (or all nodes). Because GEM uses speeds and distances in its graph, Dijkstra's algorithm effectively calculates the shortest time required to reach each node.

As mentioned earlier, Dijkstra's algorithm allows the model to assess the shortest time to reach all nodes starting from a single starting node. This could be computationally expensive if the model has multiple starting nodes to consider, such is the case if an adversary were to start at every point along the border, potentially requiring thousands of calls of this algorithm. To greatly speed up runtimes, the model inserts an additional node at zero degrees latitude and longitude (this point is arbitrary so long as it is not within a sector to process) and an edge with a speed chosen from this point to every potential starting point for the blue or red force in the scenario such that the arrival time will be what the starting time should be. When Dijkstra's algorithm begins, it will start at the newly inserted point and travel along these new edges, creating an effective initial cost to reach our real starting points. For example, if the red force is supposed to start from the border region at time = 0, the speed used along the new inserted edges will be infinite. If there are multiple blue force starting locations with different times at which they could respond, speeds will be chosen so that the starting time at each location would equate to that delay time, which may or may not be infinite. This allows all starting locations to be processed simultaneously, greatly cutting down on runtimes.

Isochrones

Following the calculation of the shortest time to reach every node in the mesh network, isochrone maps were generated to assess the area an entity can travel in a given period of time. An isochrone is defined as a contour on a

map that connects points with the same time. Isochrones are commonly used by law enforcement to determine where opposing forces can reach after a period of time. **Error! Reference source not found.** contains a representative isochrone map for an adversary incursion into the Yuma sector. In this example, an adversary is crossing the border at a single point near San Luis in a car. The map shows isochrone layers, or contours, for 10 minute intervals up to an hour. As expected, the isochrone map shows the ability of the adversary to reach further distances in a given amount of time where there are roads and highways.

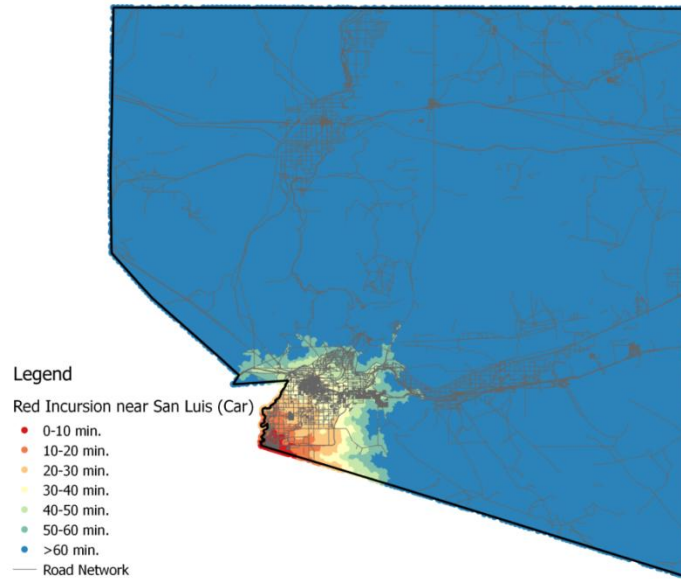


Figure 4. Isochrone map for red incursion near San Luis, Arizona

Red Force Isochrones

The red force isochrones will typically use origin points at or close to the border to represent incursions from a single location. For scenarios where the point of origin is unknown the tool can also generate isochrones from all border points simultaneously. This allows for a holistic assessment on how well the between-ports-of-entry node performs.

Blue Force Isochrones

Blue force isochrones originate at blue force response locations, such as CBP dispatch stations. Each location can have a variety of response asset types (e.g., cars, horses, etc.), each with different delays intrinsic to their mobilization or readiness level. In its current form, the tool does not factor in the presence of mobile CBP units that could be starting from an unknown or uncertain position, such as a unit on a patrol route. If the user knows that a blue force entity is in the field at a specific location, then that location and conveyance can be used.

Probability of Encounter

Once the blue and red force isochrones are generated for a given laydown of assets and incursion scenario, the probability of encounter as a function of time can be determined. The tool does this analysis using the following equations:

$$P_E(B_{t-d-s}, R_t) = \frac{B_{t-d-s} \cap R_t}{R_t} \quad (1)$$

Where R_t is the area of the red force isochrone at time t , B_{t-d} is the area of the blue force isochrone at time $t-d$, d is a time delay to account for blue force response time, and s is the amount of time it takes for a blue force asset (e.g., camera or Unattended Ground System (UGS)) to detect the red force. It may be necessary to define “win zones” or regions that, if encountered by the red force, result in an automatic victory for red or inability for law enforcement to operate within effectively, if at all. An example of a “win zone” could be an urban area or critical infrastructure. Given the existence of a “win” or exclusion zone area, W , the probability of encounter equations is modified:

$$P_E(B_{t-d-s}, R_t) = \frac{(B_{t-d-s} \cap R_t) - (B_{t-d-s} \cap R_t \cap W)}{R_t} \quad (2)$$

Blue Force Sensor Assets

The location of blue force sensor assets is highly sensitive, and as such GEM does not come with an exact force laydown structure. It does however come with notional datasets of towers and unattended ground sensors to aid the user in learning how to use the tool and how to prepare a known force laydown structure to be used with it.

Unattended Ground Sensor Emplacement and Analysis

An algorithm was generated to distribute UGS uniformly and randomly along the Northern and Southern Border within a 5 km buffer region. The number of UGS allotted to each of the sectors was also randomized, with the number of UGS along the Southern Border outnumbering those along Northern Border by 2-to-1.

Remote Video Surveillance System (RVSS) Emplacement and Analysis

Tower locations were generated algorithmically such that the towers were spaced at least 30 km apart and located on points of maximum elevation in a notional attempt to optimize their coverage. All towers were assumed to have 360 degree visibility and a height of either 50 feet or 80 feet.

Viewshed Analysis

The primary purpose of the notional tower emplacement is to demonstrate QGIS's ability to perform viewshed analysis in support of GEM. Viewshed analysis is a type of intervisibility analysis that determines the visibility of Digital Elevation Model (DEM) cells provided by the placement of towers at a height above ground level. Figure 5 contains a schematic that represents a simple viewshed analysis for a given tower over a small one-dimensional region. The area highlighted in red is visible to the tower while the area in white is not visible to the tower. QGIS performs this analysis on each tower individually and then sums up the individual viewsheds and displays them all on a single map. Figure 6 contains output for viewshed analysis of the notional placement of cameras in the Yuma sector.

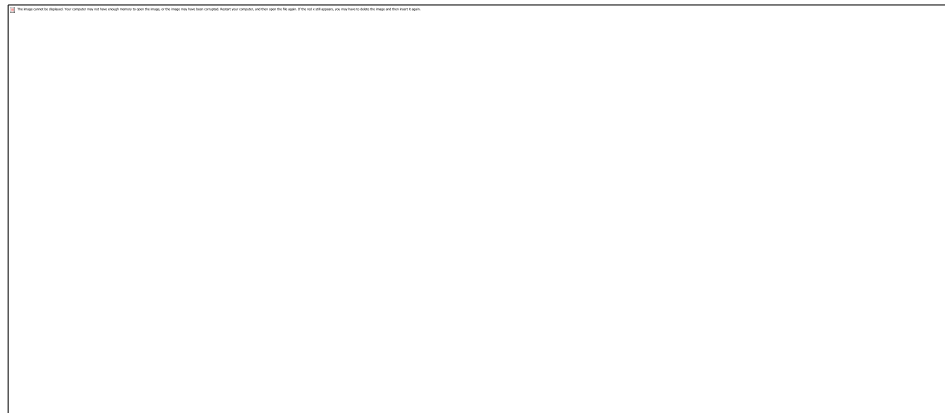


Figure 5. Representative tower viewshed

This type of analysis allows the analyst to identify visibility gaps for a given tower placement scenario and could inform the deployment of mobile assets when necessary.



Figure 6. Notional viewshed analysis for the Yuma Sector (the area is blue can be seen by the towers)

USING GEM

The GEM tool can be accessed through a plugin within QGIS. Qt Designer, an open source python based GUI building software, was utilized to develop the GEM GUI that is compatible with both Linux and Windows operating systems. The GEM GUI allows the user access to all necessary steps to complete a probability of encounter analysis; including calculation of travel graphs, placement of red and blue forces, and determination of possible encounter. The GUI has several tabs to guide a user through the analysis workflow.

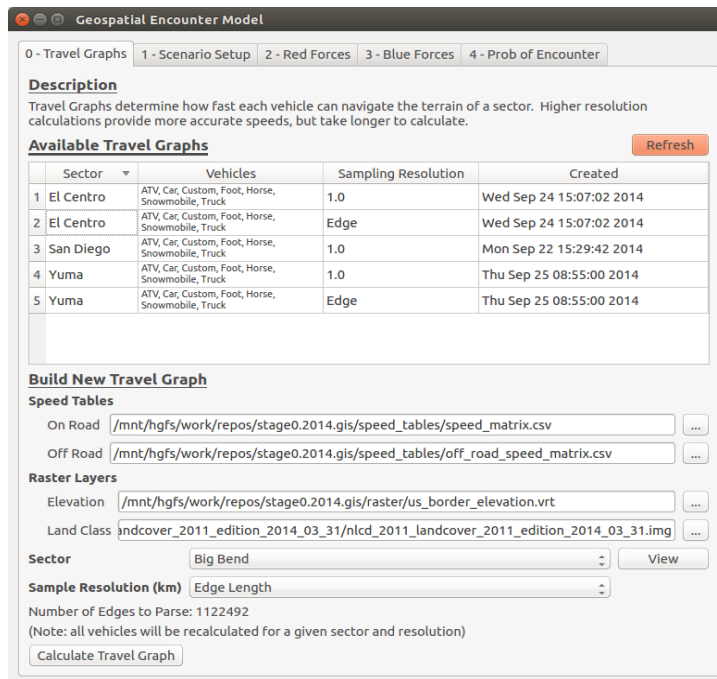


Figure 7. Travel Graph Calculation tab of the GEM GUI

The *Travel Graphs Tab* in the GEM GUI allows the user to view the available travel cost graphs (described in section 6.3 and shown in Figure 7) by sector, vehicle, and sampling resolution. All previously generated travel cost graphs will be displayed in the table. If the travel graph needed for analysis is not available, the user can calculate the appropriate travel cost graphs by selecting the appropriate speed tables, elevation file, and land class file (defaults are preloaded) and selecting the sector and resolution needed. The travel graphs will be calculated and added to the list of available travel graphs for all future scenarios.

The next step in the analysis workflow is to setup the analysis scenario within the *Scenario Setup Tab*. The user can load an existing scenario or build a new scenario by selecting the required sectors for analysis. Next the user chooses the scenario resolution, also known as the cost graph accuracy and the length of time for

the scenario to run. The scenario time parameter sets the length of time for both red and blue force travel times. After the sectors, resolution, and time are set the scenario is saved and the output path for all analysis layers is set.

Users can generate the isochrones based on the start location, scenario time, and vehicle type used by entities on the *Red Forces Tab* or *Blue Forces Tab* shown in Figure 8. The user can seed an entity or entities by map location, sector border, or input file. *Single Map Point* allows input by clicking on the map or typing a known coordinate. The sector border start location allows a Red Force to start at every point along the sector border simultaneously, simulating the locations a Red Force can reach from the border in a given time. *Input File* selection allows a user to develop a scenario involving several start locations, multiple vehicle types, and differing run times. Additionally, blue forces include a speed increase factor that allows the user to increase the travel speed over the speeds assigned in the speed tables used to generate the cost graphs. This speed factor (or multiplier) allows for the assumption that blue force vehicles may have high performance equipment and better driver training, thus allowing a blue force to travel faster than a red force. Also included are two delay parameters account for the time necessary for blue forces to respond to a red force incursion. Response delay is the time that it takes for blue forces to begin traveling after notification of a red force incursion. Sensor delay represents the time delay involved with detecting a red force movement. The sum of these two delay parameters will be subtracted from the scenario time to determine the total allowable time a blue force travels during a scenario.

After selecting the start location and vehicle type, the user chooses to create *Solid Isochrones* or *Time Step Isochrones*. *Solid Isochrones* are a single feature representing the total area a force can travel given the time and vehicle. *Time Step Isochrones* represent the same travel area, but are broken into multiple time zones to visualize the earliest a force could reach a location. Either choice in isochrone type will allow the same encounter probabilities to be computed; a choice of isochrone display is offered for user convenience. Once the *Create Isochrones* button is selected the isochrones will be generated and added to the map for display. Also, on the *Red Forces Tab* is the ability to turn on and off a few map layers, such as road networks and camera viewsheds, which may help to guide the user in the analysis.

Figure 8. Red Force & Blue Force Isochrones creation tabs of the GEM GUI

Probability of Encounter calculations are developed within the Probability of Encounter Tab (Figure 9). As a user creates red and blue force isochrones in the previous tabs, the layers are populated in the appropriate boxes in the left column. On this tab, the user can select which isochrones to use for the probability of encounter calculation by selecting the layer in the box and adding it to the box on the right. Optionally, the user may select blue force exclusion zones (also referred to as red win zones). By default the US Census Urban Zones are populated in the blue force exclusion zones and can be selected by a user, but other layers can also be added for analysis. Once the desired layers are added to the columns on the right, the user provides a name for the output analysis layer and presses the Calculate Probability of Encounter button. GEM intersects the areas that both red and blue forces can

travel, subtracts the optional excluded areas, and calculates the percentage of red travel area that blue forces can cover. The intersected area of encounter is added to the map display and the percentage is reported at the bottom of the tab.

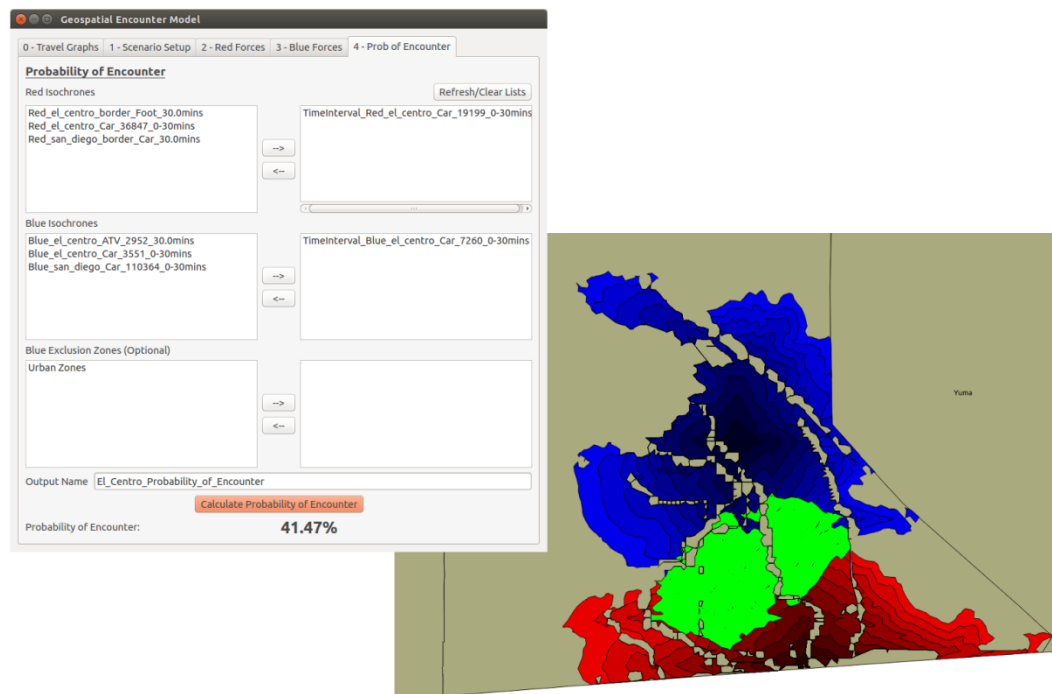



Figure 9: Probability of Encounter calculation tab of the GEM GUI and the resultant map displaying Red and Blue Force Time Interval Isochrones along with Green Probability of Encounter Area

More details on the usage of the GEM GUI can be found in the GEM GUI Manual, including annotated step-by-step directions of the entire Probability of Encounter Workflow.

GEM System

The GEM System is the combination of the GEM, geographic data covering the Northern and Southern Borders, initial analysis files including triangle mesh networks, cost graphs, and default speed tables, along with the QGIS software platform. The GEM System is delivered in a Ubuntu 14.04 virtual machine. QGIS and the GEM plugin are preinstalled with the initial analysis files and geographic data installed in the project directory. Native Windows support is available for QGIS and all code was created with interoperability in mind, but no testing has been carried out in the native Windows environment.

To utilize GEM, the user needs to first launch the Ubuntu system from within virtualization software such as VMware Player. Once the Ubuntu operating system has loaded the user will launch the QGIS environment and then the GEM plugin within. The GEM plugin can be launched from Plugins menu DNDO Modeling Plugins Geographic Encounter Model or using the  icon on the toolbar. The GEM GUI will launch providing all the necessary functionality to complete the probability of encounter calculations.

SUMMARY

Engility Corporation has developed the GEM that provides DHS with the capability to assess the potential encounter probability between law enforcement and adversaries holistically across the Northern and Southern Boundary Regions between ports of entry. This model is geospatial in nature, leveraging terrain, land type, natural boundaries, and road networks to determine how far any particular entity can get on and off-road after an arbitrary period of time. GEM utilizes the QGIS geospatial modeling and analysis software, and is provided as an easy to use plugin within this framework. Its GUI allows users to easily load their own predefined law enforcement force structure and

add locations to inject multiple adversaries into the simulation. The probabilities of encounter metrics the model generates can be used by analysts and system planners to aid in identifying vulnerabilities from potential border incursions.

ACKNOWLEDGEMENTS

This work has been supported by the US Department of Homeland Security, Domestic Nuclear Detection Office, under competitively awarded contract/IAA HSHQDC-12-A-00004. This support does not constitute an express or implied endorsement on the part of the Government.

REFERENCES

- Alden, Edward. (2013). *Measuring the Effectiveness of Border*. Washington, D.C. : Council on Foreign Relations, 2013. Before the Committee on Homeland Security and Governmental Affairs, United States Senate, 1st Session, 113th Congress.
- National Park Service. Geographic Information – National Park Service. Retrieved June 2014, from <http://www.nps.gov/gis>.
- OpenStreetMap. Retrieved June 2014, from www.openstreetmap.org.
- U.S. Customs and Border Protection, Office of Technology Innovation and Acquisition, 2011. *Northern and Southwestern Border Design Reference Mission (DRM) Version 2.0*. s.l.: U.S. Department of Homeland Security, 2011.
- United States Department of Commerce. Cartographic Boundary Shapefiles - Urban Areas. *United States Census Bureau*. Retrieved May 2014, from https://www.census.gov/geo/maps-data/data/cbf/cbf_ua.html.
- United States Geological Survey. *The National Map Viewer*. Retrieved July 2014, from <http://viewer.nationalmap.gov/viewer>.
- United States Geological Survey & United States Department of the Interior. National Land Cover Database. *Multi-Resolution Land Characteristics Consortium*. Retrieved June 2014, from <http://www.mrlc.gov/nlcd2011.php>.
- Willis, Henry H., et al. (2010). *Measuring the Effectiveness of Border Security Between Ports of Entry*. Homeland Security and Defense Center, Santa Monica: RAND Corporation.
- Shewchuk, Jonathan R. (2005). *Triangle: A Two-Dimensional Quality Mesh Generator and Delaunay Triangulator*. University of California at Berkeley. Retrieved May 2014, from <http://www.cs.cmu.edu/~quake/triangle.html>.