Analysis of Viewshed Accuracy with Variable Resolution LIDAR Digital Surface Models and Photogrammetrically-Derived Digital Elevation Models

Matthew L. Miller

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University in partial fulfillment of the requirements for the degree of

Master of Science
In
Geography

Laurence W. Carstensen, Committee Chair
James B. Campbell
Valerie Thomas

October 28, 2011
Blacksburg, Virginia

Keywords: LIDAR, viewshed, digital surface model, digital elevation model, visibility

Analysis of Viewshed Accuracy with Variable Resolution LIDAR Digital Surface Models and Photogrammetrically-Derived Digital Elevation Models

Matthew L. Miller

Abstract

The analysis of visibility between two points on the earth's terrain is a common use of GIS software. Most commercial GIS software packages include the ability to generate a *viewshed*, or a map of terrain surrounding a particular location that would be visible to an observer. Viewsheds are often generated using "bare-earth" Digital Elevation Models (DEMs) derived from the process of photogrammetry. More detailed models, known as Digital Surface Models (DSMs), are often generated using Light Detection and Ranging (LIDAR) which uses an airborne laser to scan the terrain. In addition to having greater accuracy than photogrammetric DEMs, LIDAR DSMs include surface features such as buildings and trees.

This project used a visibility algorithm to predict visibility between observer and target locations using both photogrammetric DEMs and LIDAR DSMs of varying resolution. A field survey of the locations was conducted to determine the accuracy of the visibility predictions and to gauge the extent to which the presence of surface features in the DSMs affected the accuracy. The use of different resolution terrain models allowed for the analysis of the relationship between accuracy and optimal grid size. Additionally, a series of visibility predictions were made using Monte Carlo methods to add random error to the terrain elevation to estimate the probability of a target's being visible. Finally, the LIDAR DSMs were used to determine the linear distance of terrain along the lines-of-sight between the observer and targets that were obscured by trees or bushes. A logistic regression was performed between that distance and the visibility of the target to determine the extent to which a greater amount of vegetation along the line-of-sight impacted the target's visibility.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

The analysis of visibility between two points on the Earth's surface is important to many users of Geographic Information Systems (GIS) such as military strategists, urban planners, and telecommunications companies. Most GIS software packages include the ability to compute a viewshed, the area of a terrain model visible from an observer's location. The creation of a viewshed requires the calculation of visibility between an observer's location and each cell on an elevation grid. The viewshed then portrays the terrain as being either visible or not visible from the observer's location.

Viewsheds are created using Digital Elevation Models (DEMs), digital topographic data representing the elevation of the Earth's surface. In GIS, terrain models often consist of a raster grid in which each pixel has an X and Y coordinate to determine its location in a geographic coordinate system, as well as a Z attribute to represent the elevation of that pixel. DEMs are often known as "bare-earth" elevation models because they represent only the topography of the Earth's terrain while Digital Surface Models (DSMs) include surface features such as buildings and trees (Maune et al., 2007).

Most DEMs have been created using photogrammetry, the process of determining terrain topography using aerial photographs. Photogrammetry relies upon the use of *stereoscopic parallax* using two aerial photographs depicting the terrain from slightly different positions. The two perspectives enable the viewing of the terrain in three dimensions, allowing for perception of relief in the terrain and the determination of the terrain's elevation (Paine & Kiser, 2003). A series of *masspoints* are created representing the elevation of the terrain at variously spaced

locations in the aerial photographs. Additionally, breaklines may also be created from the photographs to represent sudden changes in the topography such as cliffs or ridge lines.

A newer technology used to create DEMs and DSMs is Light Detection and Ranging (LIDAR) which uses a pulsating laser attached to an airplane to sweep across the terrain. The distance from the plane to the surface can be calculated using the time delay between the pulse's emission and its return to the sensor. Measurements by laser are more accurate than those of photogrammetry and therefore create a more detailed terrain model. The pulse rate of the laser results in more sample points than photogrammetry, also contributing to greater precision. In addition, LIDAR pulses reflect off surface features such as trees and buildings, making LIDAR more suitable for the creation of detailed DSMs.

1.1 Objective

This project analyzes several questions raised in previous literature about viewsheds that have not been fully studied and quantified. The primary goal is to compare the accuracy of visibility based on both photogrammetrically-derived bare-earth DEMs and LIDAR-derived DSMs. A visibility algorithm will be executed on both types of terrain models between a series of observation points and target points. Because a DSM generated with LIDAR contains a highly detailed representation of vegetation and man-made structures, this project will measure how often those surface features cause a different result from visibility analysis based on a bare-earth model. The differences in the analysis for the DEMs and DSMs will then be compared to the real-world results of a field survey.

Additionally, the project will analyze the effect of resolution on both the DEMs and DSMs to determine whether there exists an optimal resolution for viewshed accuracy. Past

research suggests that DEMs and DSMs are most accurate when the grid resolution matches the average spacing of the sample points in the dataset from which the DEM or DSM was created (Li, 1992; 1994; Gong et al, 2000). Decreasing or increasing the resolution beyond the average point spacing may decrease the accuracy of the terrain model and likewise decrease the accuracy of the viewshed. To test this hypothesis, a series of DEMs and DSMs will be created at varying resolutions from photogrammetric and LIDAR datasets to determine the degree to which changes in resolution affect the accuracy of visibility predictions.

This project will also conduct a visibility analysis with all vegetated or tree-lined areas of the LIDAR DSMs delineated so that any blockages of visibility due to vegetation may be removed from consideration. Because the basic visibility algorithm considers vegetation in a DSM to be entirely opaque, eliminating obstructions of visibility that occur in vegetated areas will produce an estimate of how often vegetation determines whether a target is visible or not. The analysis on the LIDAR DSMs that do not include vegetation will also serve as a comparison to the photogrammetry-based DEMs which do not include any surface features.

Finally, an analysis will be performed on the LIDAR surfaces using only the targets which are blocked from view by vegetation. The linear distance that the observer's line-of-sight travels through vegetation will be compared to the degree to which the target is visible in the field. An ordinal regression between the two should reveal the relationship between the amount of vegetation along the line-of-sight and the target's visibility in the real world. One would naturally expect a line-of-sight that passes through a greater distance of vegetation would make the target less visible. An ordinal regression between the two factors should provide a means to estimate the likelihood of a target's being visible, blocked, or partially visible based on the amount of vegetation through which the line-of-sight passes in the LIDAR DSMs.

Chapter 2: Literature Review

2.1 Viewsheds and Visibility – Concepts and Use

The concept of terrain visibility addresses whether a target point on the Earth's surface is visible from an observer's geographic location. The user seeks to answer a simple question: Is Point A (the target's location) visible from Point B (the observer's location)? Whether a single target point is visible from an observer's location is known as *intervisibility*, while the set of all visible points on the terrain surrounding the observer's location is known as a *viewshed.* ESRI, the software company that develops the popular ArcGIS software, states that calculating a viewshed is "useful when you want to know how visible objects might be - for example, from which locations on the landscape will the water towers be visible if they are placed in a particular location, or what will the view be from a road?" (ESRI ArcGIS Resource Center, 2010a).

Most GIS software packages include the ability to calculate both intervisibility and a viewshed. ESRI's ArcGIS, one of the most commonly used GIS software packages, includes functions for analyzing both point-to-point visibility as well as determining the visibility of all points in a terrain model from an observer's location. To determine intervisibility between the observer and target, the software analyzes the terrain along the line-of-sight (LOS) between the observer's location and the target's location. The software then compares the elevation vector between the observer and target with the elevation of the terrain along the LOS. If the terrain along the LOS is elevated so that it blocks the view of the target, then the target's location is determined to be not visible. Otherwise, if no terrain along the LOS blocks the view, then the target's location is deemed to be visible. The creation of a viewshed simply extends the application of the intervisibility calculation to the entire terrain surrounding the observer; the

intervisibility algorithm is performed on all surrounding terrain and results in a map illustrating which locations are visible to the observer and which are not.

The desire to determine terrain visibility has great importance to a variety of GIS applications. It is commonly used by military strategists seeking to conceal the movement of soldiers, urban planners seeking to predict the views of residential developments, and telecommunications companies seeking to predict the best location to place transceivers, such as cell phone towers. Terrain visibility has had an enormous impact on military history and has determined the outcomes of battles by either concealing or revealing troop movements. Ehlen and Abrahart (2004) used GIS to analyze the role of terrain visibility in the Civil War battle of Perryville, Kentucky, where terrain relief led to misjudgments of the distance of enemy forces, resulting in inaccurate targeting of cannon fire, while Confederate forces launched successful infantry attacks on Union positions by approaching through small valleys that concealed their advance. Ehlen and Abrahart (2002) also describe how the visibility of soldiers to enemy artillery rendered certain areas of low-lying terrain impassable during the Battle of Fredericksburg, Virginia in 1862.

Visibility is also a prime factor in the siting and construction of fortifications and defensive positions. Guth (2004) describes the development of fort construction in the 17th through 19th centuries in which forts were designed in a star-like shape to maximize the visibility of approaching enemy forces and to maximize the fort's weapons fan, the field in which cannon fire could effectively target approaching forces. Entrenchments during the First World War were also constructed in jagged-line patterns along the top of ridgelines to give defending troops the maximum visibility of the ground below (Doyle et al., 2002). Major engagements between military forces often occur over strategic high ground which gives the controlling force superior

visibility over the terrain, such as the World War II battle of Monte Cassino, Italy, which was fought to control a hill overlooking a strategic mountain pass between Naples and Rome (Ciciarelli, 2002).

Viewsheds and visibility analyses also find common use in the field of landscape planning to create scenic vistas for commercial real estate developments (Fisher, 1995; 1996). Housing developments that promise scenic views and tourist resorts that overlook beaches or mountain vistas seek to offer the best views to their clients, making viewshed analysis a significant factor in the choice of site location. Bishop (2002) describes the use of viewshed analysis to determine the visual impact of a wind turbine farm on the views of surrounding residential neighborhoods in the countryside of England. Bishop and Gimblett (2000) also used visibility analysis as a factor in creating a GIS model to predict the popularity of recreational areas among tourists, finding that a 3D model illustrating scenic views can help determine which hiking and biking trails will be most popular.

Another application of visibility analysis seeks to determine a variation of a viewshed known as a *commshed.* Telecommunications companies want to find locations for cell phone towers and similar transceivers that provide maximum signal reception to the largest number of people in a geographic area as possible. A commshed, therefore, describes the area in which a telecommunications device such as a cell phone can receive a signal from a transceiver, such as a cell phone tower. As in visibility analysis, a telecommunications signal may be blocked by terrain features such as valleys and ridges. Previous research has analyzed the effect of terrain on the reception of radio wave transmitters that rely on line-of-sight visibility for quality signal reception (Rose, 2001; Cash, 2003). The ability to use GIS viewshed functions to predict a commshed was analyzed by Dodd (2001), who compared viewsheds created with ESRI ArcGIS

to fieldwork involving the broadcast and receipt of a high-frequency line-of-sight radio wave transmission.

2.2 The Determination of Visibility

Before the advent of GIS, cartographers relied on topographic maps to predict terrain visibility by drawing a series of radiating transects from the observer's location. Topographic profiles would be created along each transect which would be used to determine which areas of ground were not visible to the observer (Yoeli, 1985). The process was very time-consuming and depended on the accuracy and detail of the topographic base map. The spacing of the topographic contours, as well as the number of radiating transects used, had a direct impact on the reliability of the resulting visibility map.

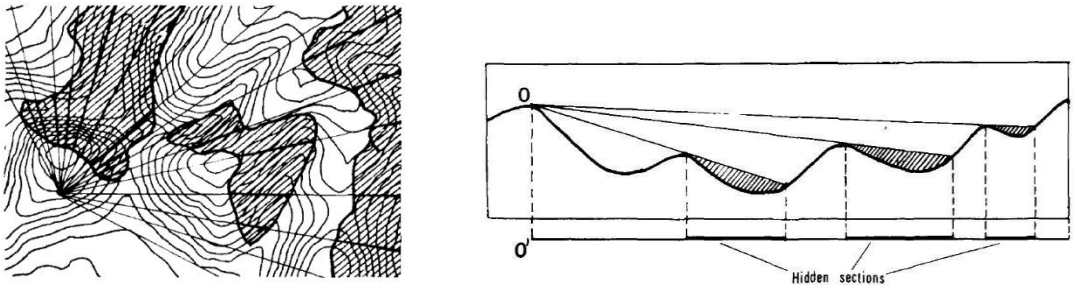

Figure 1 - Early intervisibility maps were created by drawing a set of radiating transects across a topographic map from the observer's location. A terrain profile (right) shows areas blocked from an observer's location (Point O). Based on an analysis of the terrain profile along each transect, a map could be created showing areas not visible to the observer (left). (Yoeli, 1985; Reprinted by permission of University of Toronto Press Incorporated).

The accuracy of the topographic profile along the observer-to-target line-of-sight is crucial for an accurate prediction of visibility. If the target is behind a hill or ridge that blocks the LOS, the observer's view of the target will most likely be blocked. Depending on the

elevation of terrain along the LOS, areas just behind a hill or ridge may be blocked from sight, while more distant areas remain visible. In the following illustration, the observer has a clear line-of-sight to the ground on the right-most side of the terrain profile because the LOS is constantly above the intermediate terrain:

Image courtesy of L.W. Carstensen

If the observer seeks to view terrain that falls just behind the central hill, however, his line-of-sight will be blocked by the hill. Thus, a series of lines-of-sight can be constructed to illustrate which sections of terrain are visible and which are not:

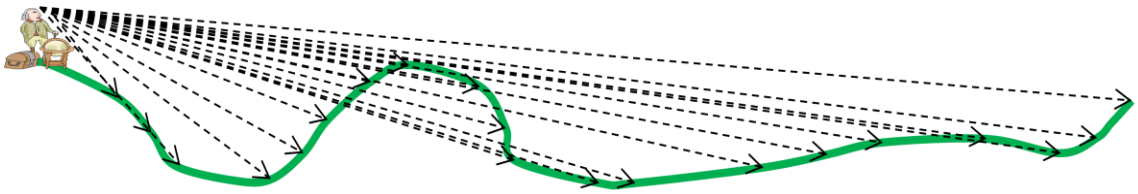

Image courtesy of L.W. Carstensen

Any line-of-sight that crosses the terrain profile indicates that its end point is not visible to the observer. Therefore, the valley just below the observer, as well as the left side of the hill, would be visible, but the backside of the hill and the valley behind it are not visible. The upward slope on the far right edge of the terrain, however, would be visible since the line-of-sight to that terrain remains above the center hill.

2.3 The Development of Viewsheds in GIS

The earliest digital visibility maps were created using raster grids to represent the terrain's elevation. Yoeli (1985) described the creation of a program that analyzed a grid of elevation values surrounding an observer's location. The program determined the visibility of a target grid cell by creating a terrain profile from the grid cells along the line-of-sight between the observer and the target cell. Yoeli's program repeated the process for each cell in the grid, assigning a value of one to cells determined to be visible and zero to those determined to be blocked from view. The resulting grid, therefore, predicted which areas of the surrounding terrain would be visible to the observer and which were not.

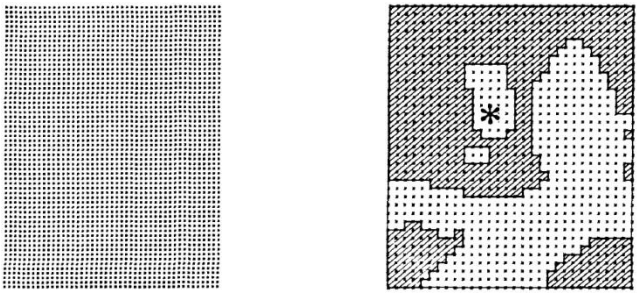

Figure 2 - Yoeli's 1985 viewshed program relied on a grid of elevation points to represent the terrain (left). Grid points determined to be blocked to the observer (marked by the asterisk at right) are hashed while visible points are not (Yoeli, 1985; Reprinted by permission of University of Toronto Press Incorporated).

The mathematical determination of whether the line-of-sight crosses a section of intervening terrain relies on the comparison of the angle of the observer-to-target LOS with the angle of the LOS to each intermediary point on the terrain, as described by Yoeli in 1985. Most intervisibility algorithms begin by determining the tangent of the angle created by the difference in elevation of the observer and the target. The algorithm then proceeds to calculate the tangent

of the angle between the observer and each point on the terrain model that falls along the observer-to-target LOS. If the observer is lower than the target, the view is considered blocked anytime the tangent of the angle between the observer and an intermediate grid cell is greater than the tangent of the observer-to-target angle. If the observer is higher than the target, however, the view is considered blocked anytime the tangent of the angle between the observer and an intermediate grid cell is less than the tangent of the observer-to-target angle.

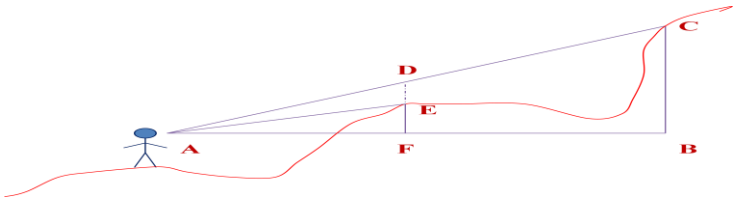

Figure 3 - Determination of visibility by comparing tangents of sightlines (Image courtesy of L.W. Carstensen).

The comparison of angles is illustrated in Figure 3 above. The LOS between the observer's eye level (point A) and the target (point C) is continually above the terrain in between. A visibility algorithm determines this by comparing the tangent of angle DAF to the tangent of each angle formed by each point on the terrain between A and C. For the hill in the middle of the terrain profile, the highest elevation lies at point E. When the tangent of angle DAF is compared to the tangent of angle EAF, it is clear that the tangent of EAF is smaller. Therefore, point E does not block the LOS between the observer and the target.

2.4 How Terrain Profiles are Generated from a Raster Grid

GIS software commonly represents terrain elevation as a raster grid in which each grid cell contains an elevation value, usually measured in feet or meters above the vertical datum used for the terrain model. Each grid cell is denoted by its column and row number, which corresponds to its X and Y coordinates in the geographic coordinate system. With the elevation value, or Z attribute, each grid cell has an X, Y, Z coordinate location in three dimensional space. The width of the grid cells denotes the raster's resolution. Higher resolution rasters (with smaller grid cells) yield more accurate representations of the terrain.

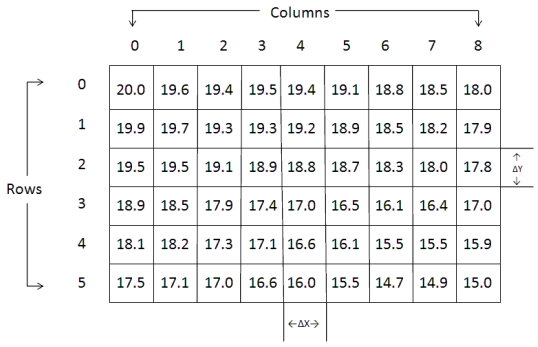

Figure 4 - An example of a raster grid used to represent topography. Each grid cell has its own elevation value and its X and Y coordinates represented by column and row numbers. The resolution of the grid is determined by the width and height of the cells (ΔX, ΔY).

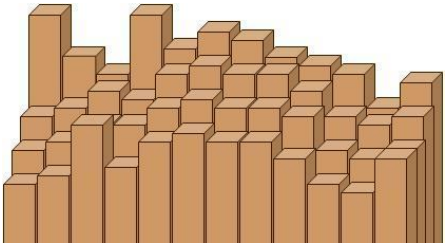

Figure 5 - A raster grid viewed three dimensionally as a series of columns whose heights reflect the elevation of each grid cell (Image courtesy of L.W. Carstensen).

When determining intervisibility, the observer's location is considered to be the grid cell that contains the geographic location of the observer. The target's location, likewise, is the grid cell that contains its geographic location. The elevation of both target and observer are the elevation values of their respective grid cells. The visibility of the target from the observer, therefore, depends on whether the LOS is blocked due to the elevation of any intermediate grid cells.

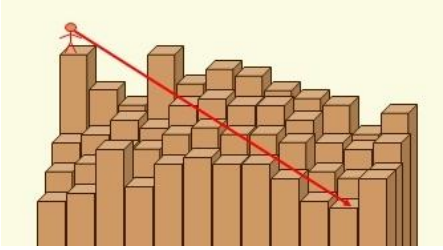

Figure 6 - The line-of-sight from observer to target on a raster grid elevation model (Image courtesy of L.W. Carstensen).

To create a terrain profile for the calculation of intervisibility, the software must select all grid cells along the line-of-sight from the observer and the target. When the observer and target are located in the same row or column of the raster grid, this process is very simple; the terrain profile in this case consists of the elevation values of the intermediate grid cells along that particular row or column. When the observer and target are located in different rows and columns, however, the process becomes more complex. A commonly used algorithm to determine elevation values along the LOS in this case is the *Bresenham algorithm*, which determines which grid cells lay along the LOS between the observer and target. The algorithm was first described by J.E. Bresenham (1965) to control digital plotters, but is also effective for determining which cells in a raster grid fall along the LOS.

The Bresenham algorithm operates by considering the difference in the X and Y coordinates of the observer and target. Since the observer and target are both located in a grid cell, the X coordinate of each corresponds to their respective column numbers and the Y coordinate of each corresponds to their row numbers. By calculating the row difference and column difference between the observer and target, the algorithm selects each cell lying along the line directly between them.

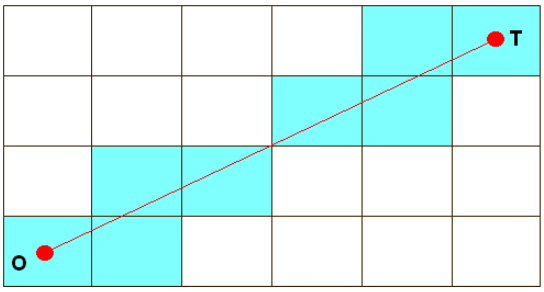

Figure 7 - The Bresenham algorithm is used to determine which grid cells are crossed by the line-of-sight between two points. In this example, the observer (O) and target (T) are separated by three rows and five columns. The cells shaded blue are those determined by the algorithm to fall along the line-of-sight.

Once the grid cells along the line-of-sight have been determined, the elevation value of each grid cell is loaded into an array, which holds the elevation values of the terrain profile between the two points. However, the LOS does not cross each cell at its center, with the exception of the beginning and end cells. Therefore, the terrain profile may be further refined by interpolating the elevation value at the approximate location at which the LOS enters and leaves each grid cell as illustrated in the following diagram:

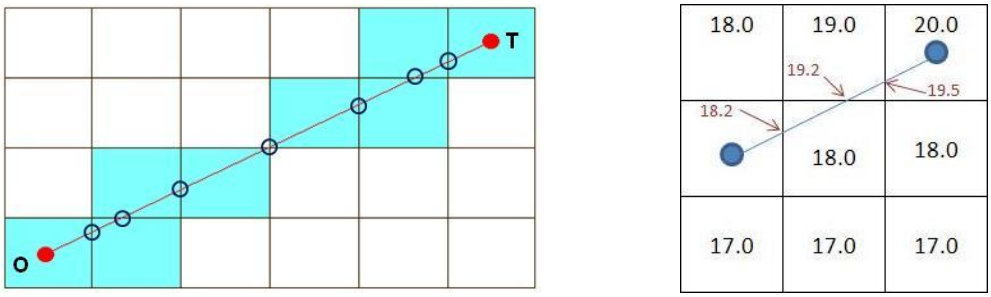

Figure 8 - Interpolation may be performed to determine the elevation values at each point where the LOS enters and leaves each grid cell, as marked by the circles (left). When an LOS intercepts the boundary of a grid cell, the terrain elevation at that point may be approximated using the values of the neighboring cells (right).

The elevation values at each point where the line-of-sight enters and leaves a cell are determined by considering the elevation values of the adjacent cells. The elevation is then interpolated using the adjacent elevation values. For a linear interpolation example, if the LOS enters a cell 3/10 of the cell's width on its left side, then the elevation value at that point is 3/10 of the difference between the cells immediately above and below. Likewise, if the LOS enters a cell from the top or bottom sides, then the elevation value at that point is interpolated using the values of the cells on the right and left. The interpolation may then be repeated when the LOS exits the cell on any side. Figure 8, at right, shows approximations of elevation values where the LOS crosses cell boundaries based on the elevation values of the adjacent cells.

After the cells underneath the LOS are selected and the elevation values are interpolated, these values create the terrain profile needed for determining visibility of the target from the observer. The intervisibility algorithm based on the concepts outlined by Yoeli (1985) and Fisher (1991) is performed on the elevation values to determine whether the intermediary terrain blocks the view of the target from the observer's location. The use of a grid-based DEM, the

Bresenham algorithm, and the interpolation of all points at which the LOS intersects a grid cell boundary have been further described by Fisher (1992; 1994).

When the combination of the Bresenham algorithm and the intervisibility algorithm are used iteratively to determine visibility from the observer to all cells in the grid, a viewshed is created (Figures 9 & 10):

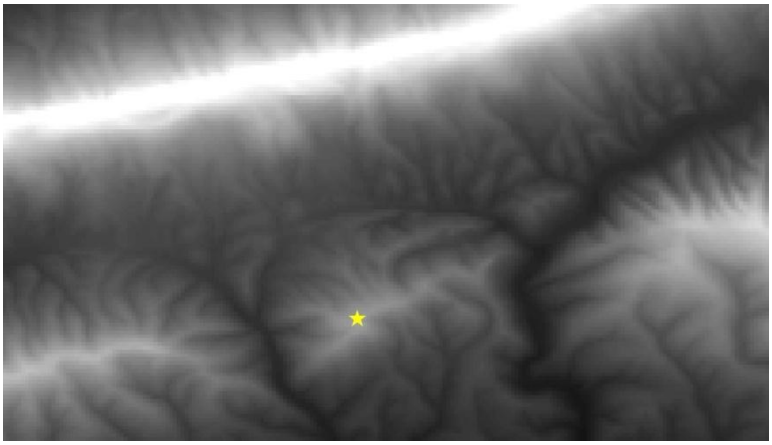

Figure 9 - A Digital Elevation Model created from the US Geological Survey's National Elevation Dataset of the Blue Ridge Mountains near Blacksburg, Virginia. A hypothetical observer's location is marked with the yellow star (Image by author using ArcGIS).

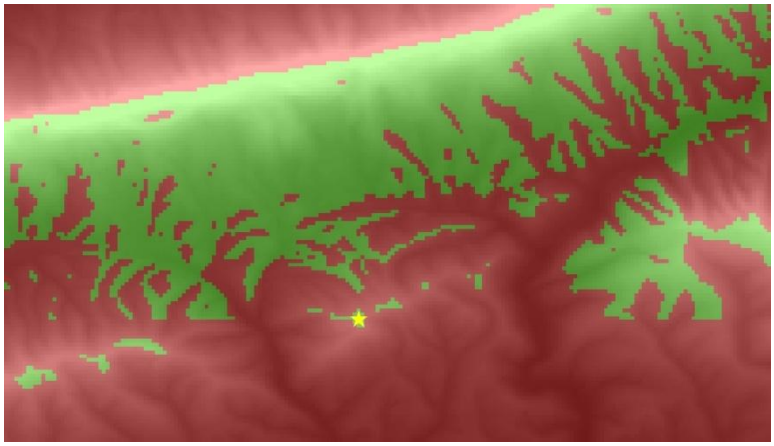

Figure 10 - A viewshed created from the USGS DEM. The areas in green are predicted to be visible to the observer and those in red are not. The viewshed was created with ESRI's ArcGIS software (Image by author using ArcGIS Spatial Analyst).

2.5 The Creation of Digital Elevation Models in GIS

Because the determination of visibility requires a raster grid of terrain values, the reliability of any viewshed depends upon the quality of the terrain model upon which it is based. A series of factors, however, affects the quality of a terrain model. These include the method used to acquire elevation data, the interpolation technique used to convert the data into a raster grid, and whether or not the model includes features such as buildings or vegetation. A great deal of previous research has been conducted on each of these factors. Any factor that influences the accuracy of the digital terrain model will, in turn, affect the reliability of visibility predictions.

When dealing with terrain models one must distinguish the type of terrain model one is working with. Raster grids of elevation values are often referred to with the terms Digital Terrain Model (DTM), Digital Elevation Model (DEM), or Digital Surface Model (DSM). While some GIS users may use these terms interchangeably, others draw distinctions among the types. Maune, et al. (2007) explains that a DEM consists simply of terrain elevation values "void of vegetation or manmade features," and are often referred to as "bare-earth" models. A DSM, on the other hand, does include vegetation and man-made features such as buildings, roads, or bridges, which can create a significant difference between its elevation values and those of a "bare-earth" DEM. Maune et al. (2007) further distinguish between a DTM and a DEM by defining a DTM as a bare-earth model that includes additional, distinctive topographic features such as cliffs or ridgelines. These features are represented by breaklines which add additional detail to a terrain model.

Since the accuracy of the data determines the quality of the model, the technique used to acquire the model's data constitutes a significant factor in the model's accuracy. Traditionally,

most elevation models have been based on photogrammetry, which uses two overlapping aerial photos, taken from different positions, to create a three-dimensional stereoscopic image of the terrain. An analyst can then calculate the elevation, or *Z-coordinate*, of a point, as well as its *X* and *Y* coordinates in a chosen geographic coordinate system. A series of elevation points, known as masspoints, as well as breaklines that represent significant features such as cliffs or ridgelines, can then be used to generate a DEM or DTM (Molander, 2007).

A more recent development in the acquisition of terrain data uses an airborne laser to scan the earth's surface. Known as *Light Detection and Ranging* (LIDAR), an airplane with a mounted laser flies across the terrain while the laser sends a series of pulses that strike the ground and reflect back to a sensor. Since the speed of light is constant, the time required for each pulse to strike the ground and return to the sensor can be used to calculate the distance between the ground and the aircraft. This calculation is combined with data from a Global Positioning System which records the aircraft's geographic location and altitude. Additionally, the plane carries an Inertial Measurement Unit (IMU) which records its roll, pitch, and yaw, and thereby constantly monitors the plane's orientation in three-dimensional space.

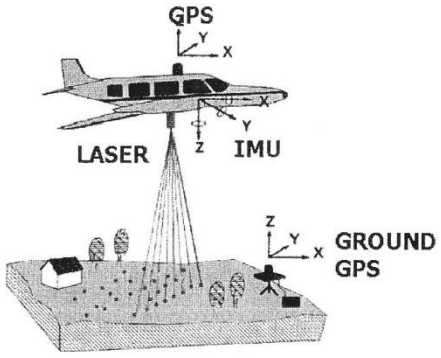

Figure 11 - LIDAR data is obtained with an airborne laser that scans the ground underneath. A GPS unit measures the sensor's location and altitude while an Inertial Measurement Unit (IMU) tracks the roll, pitch, and heading of the aircraft. (Reprinted with permission from the American Society for Photogrammetry and Remote Sensing.)

A LIDAR sensor generates a high frequency of laser pulses, as much as 250 kHz (250,000 laser pulses per second), enabling the creation of a very dense dataset of elevation readings (Lemmens, 2007). Even lesser frequency systems, such as the more commonly used 5-10 kHz systems, generate extremely detailed models of the earth's surface (Fowler, 2007). LIDAR also produces elevation readings of much greater accuracy than those derived through photogrammetry, with most LIDAR datasets having a Root Mean Square Error of 10-15 cm at a 95% confidence level (Fowler, 2007).

After elevation data is collected, a raster grid must be generated. The raw data for both LIDAR and photogrammetry, however, consists of unevenly spaced point elevation values and therefore it must be *interpolated* in order to create a raster grid. According to ESRI, "interpolation predicts values in cells for a raster from a limited number of sample data points (ESRI ArcGIS Resource Center, 2010b)" The goal of interpolation, therefore, is to create an evenly spaced set of elevation values that represent the terrain's topography from a limited and

often unevenly spaced set of samples. The following illustration shows how a raster grid is derived through interpolation:

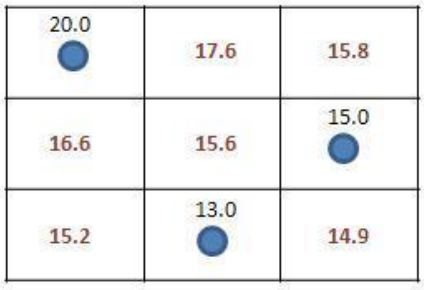

Figure 12 - Raster grid interpolation from a limited set of sample points. The interpolation of the points (blue) uses the Inverse Distance Weighted (IDW) method which approximates the values of each cell in the grid based on the values of the sample points.

Interpolation relies on the principle of spatial autocorrelation, in which points near each other are more likely to correlate than points farther away. In the illustration above, the center grid value is determined to be 15.6 because it lies between the sample points of 20.0, 15.0, and 13.0. The prediction for the center cell's value is most heavily weighted by its neighboring cells, 15.0 and 13.0, and less heavily weighted by the 20.0 value on its upper left. Likewise, the upper-center and upper-right cells are determined to be 17.6 and 15.8, respectively, with the upper-center cell having more weight from the 20.0 value in the upper left cell and the upper-right cell having more weight from the 15.0 cell just below it. Interpolation to a raster, therefore, allows a raster grid representing the earth's surface to be derived from the limited set of elevation points generated by LIDAR and photogrammetry.

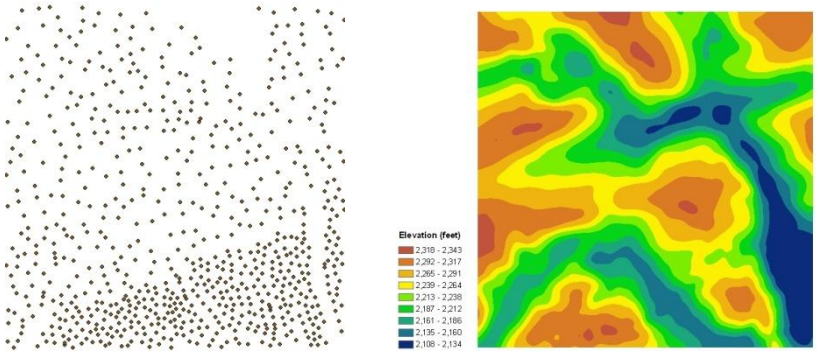

Figure 13 - Interpolation allows for a scattered sample of elevation measurements (left) to be used to create a raster grid representing the terrain topography (right). (Sample data from the Virginia Base Mapping Program near Wytheville, Virginia; Image by author using ArcGIS).

Because interpolation involves approximating nearby values from the sample set, there exist a number of different mathematical techniques for interpolation. The most commonly used techniques are known as Inverse Distance Weighted (IDW), Natural Neighbor, Spline, and Kriging. Each technique has its own advantages and disadvantages and GIS analysts often have differing preferences for each depending on the type of data being interpolated. According to Lam (1983, p. 130), "none of them is superior to all others for all applications, and the selection of an appropriate interpolation model depends largely on the type of data, the degree of accuracy desired, and the amount of computational effort afforded."

The simplest interpolation technique is Inverse Distance Weighted, which predicts the value of a grid cell by linearly weighting scattered sample points. According to Amidror (2002, p. 166), each interpolated point "is a weighted average of the values of the scatter points, and the weight assigned to each scatter point diminishes as the distance from the interpolation point to the scatter point increases." In other words, the weight assigned to a sample point is the inverse of that sample point's distance from the grid cell being interpolated. Those sample points that are closest to the grid cell being approximated will therefore have the greatest influence, while

those further away will have proportionately less influence. While IDW works well for a dense set of sample points, its accuracy may be diminished if used on datasets that are uneven or contain too few sample points (Liu, 2008). In addition, IDW may suffer from a dimpling effect in which small pits are formed near the location of the sample points (Childs, 2004).

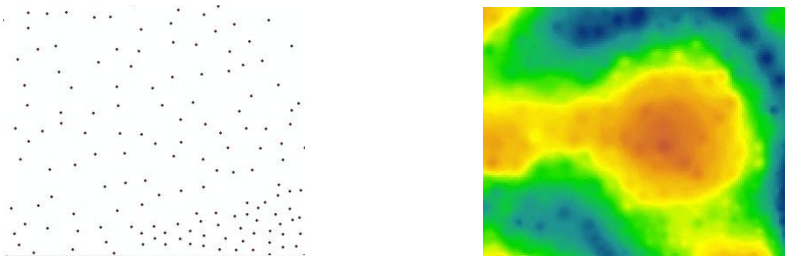

Figure 14 - A set of sample elevation points (left) interpolated with the IDW function (right). (Sample data from the Virginia Base Mapping Program for Wytheville, Virginia; Image by author using ArcGIS).

The Natural Neighbor technique relies on interpolating a point based on its most appropriate neighboring points. It creates a set of Voronoi, or Thiessen, polygons around the sample points. Each Voronoi polygon contains the area around a sample point that is closer to that point than any other sample point (Amidror, 2002). For each point being interpolated, a new Voronoi polygon is created and the value at the point is determined by the proportion with which its Voronoi polygon overlaps the Voronoi polygons of the sample points (Amidror, 2002). Natural Neighbor is useful for maintaining smoothness of the interpolated surface, but will not exceed the range of the input data, making it less suited for sparse or unevenly spaced sample sets (Maune, 2007).

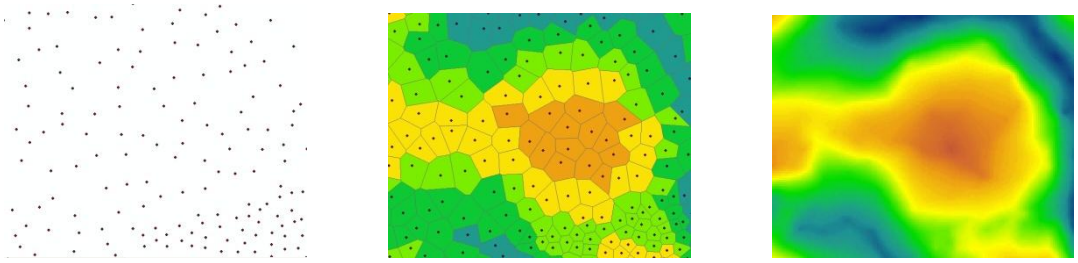

Figure 15 – The set of sample elevation points from Figure 14 (left) is grouped into a set of Voronoi polygons (center). When interpolating with Natural Neighbor, the value of each interpolated point is weighted by the proportion of that point's Voronoi polygon that overlaps the Voronoi polygons of the sample points. The final interpolation is shown at right. (Image by author using ArcGIS).

The spline technique seeks to create a smooth surface that passes through all sample points while minimizing its overall curvature (Childs, 2004). Spline can excel over IDW when predicting the presence of ridges and valleys that may not be adequately sampled. Denser datasets will also generate smoother surfaces with the spline technique but the drawback, however, will be increased computational time (Childs, 2004).

Kriging is a more complex form of an interpolation method based on geostatistical analysis. It not only considers the distance of the sample points from the point being interpolated, but the degree of variance in the sample data (Liu, 2008). The variation in the surface values is assumed to be correlated with the distance and direction from the point of interpolation (Childs, 2004). Kriging is less commonly used in the creation of DEMs and DTMs, however, and is most often used in the fields of geology or soil science (Liu, 2008).

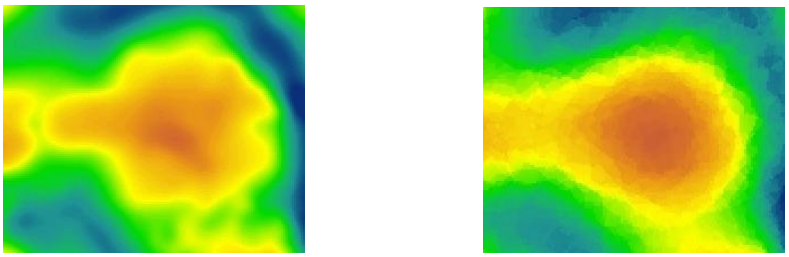

Figure 16 - The set of elevation points from Figure 14 are interpolated with the Spline method (left) and the Kriging method (right). (Image by author using ArcGIS).

Accuracy of scattered point interpolation may be further increased with the use of breaklines which trace significant topographic features on the terrain such as cliffs, ridges, or bodies of water. Breaklines are often generated using stereoscopic photogrammetry and included with masspoints in photogrammetrically generated elevation datasets. Breaklines may be used as barriers in the interpolation process to prevent sample points on the opposite side of the barrier to be used in the interpolation of a grid cell (Childs, 2004). By excluding points on opposite sides of a breakline from consideration in interpolation, a model may be created that more accurately represents terrain with significant changes in elevation (Li, 1992; Gong et al., 2000). The following diagram illustrates the difference when interpolating with breaklines:

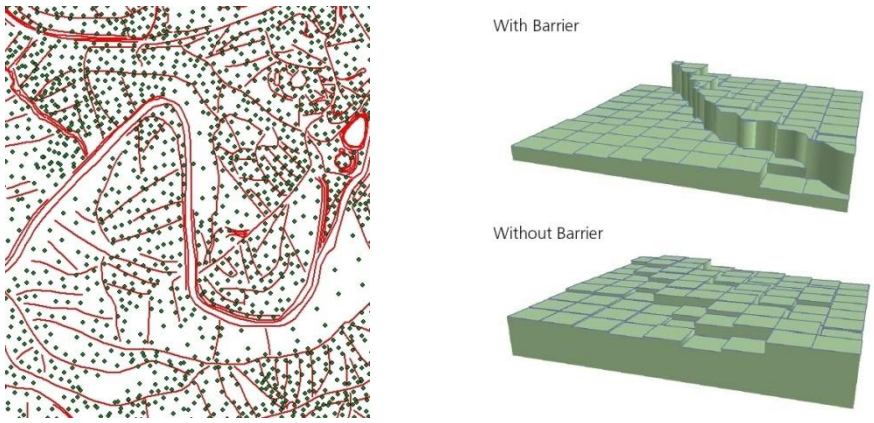

Figure 17 - Breaklines (left, in red) may be used as barriers in the interpolation process to create a more accurate representation of topographic features such as cliffs, ridges, or streams. A comparison of the resulting surface with barriers and without is shown at right. (Left image by author, from Virginia Base Mapping Program data for Wytheville, Virginia; right image courtesy of Childs, 2004).

2.6 Issues Affecting the Accuracy of Viewsheds

Because most viewsheds in GIS rely on raster grids generated from sampled terrain data, the accuracy of the input data plays a crucial role in determining the accuracy of the viewshed. A terrain model based on less accurate elevation data will obviously yield less accurate visibility predictions and viewsheds. The difference in accuracy in terrain models generated from photogrammetry compared to those created from LIDAR, therefore, will likewise result in viewsheds of varying accuracies.

The most significant source of error in the creation of a viewshed is the error in the original dataset used to create the raster grid upon which the viewshed is based (Ruiz, 1995; Riggs & Dean, 2007). All elevation datasets, whether generated from photogrammetric methods, LIDAR, or field surveying will have a certain margin of error, measured by the Root Mean Square Error (RMSE), and a statistical confidence level, often 90 or 95%. The true elevation of all sample points, therefore, will be within the RMSE of the stated value 90% or 95% of the

time, depending on the dataset's confidence level. When a raster based terrain model is generated from a set of sample points, the error in that dataset will transfer into the DEM. A point dataset with greater error generates a DEM with greater error and, in turn, any viewsheds created from that DEM will be less accurate (Ruiz, 1995).

The amount of error in a sample dataset often depends on the type of terrain it represents. Fisher (1998) points out that the consideration of a dataset's RMSE makes the faulty assumption that the error is evenly distributed throughout the dataset. Indeed, areas with higher terrain relief, such as hilly or mountainous areas, have been shown to have a higher degree of error in most elevation datasets which results in less accurate raster-based DEMs (Li, 1992; Gong et al., 2000). Ruiz (1997) likewise found that rougher terrain generates significantly less reliable viewsheds.

Another factor influencing the accuracy of DEMs is the sampling interval at which the elevation data is measured. A smaller number of sample points within a given area will produce a less accurate DEM while a higher density of points generates a more accurate DEM (Li, 1992). Gong et al. (2000) found that the degree of error in a DEM decreases linearly with the increasing density of the sample points. A higher density of sample points, therefore, generates a more accurate terrain model.

While early studies of error in DEM generation focused on photogrammetrically-derived datasets, LIDAR-derived datasets have been shown to be influenced by the same factors when evaluating their accuracy. Hodgson & Bresnahan (2004) and Su & Bork (2006) both found a relationship between terrain roughness and accuracy of LIDAR datasets. Hodgson et al. (2003) found that differing types of vegetation may impact the vertical accuracy of LIDAR. However, LIDAR's accuracy appears to be most significantly affected by the *steepness* of the terrain, making terrain with a high slope, such as the walls of a valley, the least accurate elements of a

LIDAR terrain model (Huising & Gomes Pereria, 1998; Hodgson & Bresnahan, 2004; Bater & Coops, 2009). The altitude of the airplane has also been found to affect the accuracy of a LIDAR-derived terrain model with lower altitude flights yielding more accurate data (Garcia-Quijano et al., 2008). Hodgson & Bresnahan (2004) also point out that varying flying height is one factor in decreasing accuracy over rough terrain since the altitude of the aircraft will vary relative to ridges and valleys. Additionally, a LIDAR dataset with fewer sample points, as is the case when a lower-frequency laser is used, will lead to a less accurate terrain model (Fowler, 2007).

2.7 Error in Interpolation

In addition to the quality of the initial dataset used to generate a DEM, the interpolation technique used to create a raster-based terrain model from a scattered set of sample points may introduce error. A great deal of research has been conducted examining the superiority of various interpolation methods. While Lam (1983) noted that no interpolation technique is entirely superior to all others, certain techniques may be preferred depending on the type of terrain or the characteristics of the dataset. The use of different interpolation techniques may therefore be a significant factor when considering sources of error from DEMs derived from photogrammetric datasets as opposed to those generated using LIDAR.

When interpolating photogrammetric datasets, the use of breaklines has been found to significantly enhance the accuracy of the DEM, especially in rough terrain (Li, 1992; Li, 1994; Gong et al., 2000). Not all interpolation techniques allow the use of breaklines, however, and in ESRI ArcGIS, only IDW and Kriging support their use (Childs, 2004). Therefore, if breaklines are available in a photogrammetric dataset, using them with either IDW or Kriging would likely

create a more accurate DEM than an interpolation method that does not support breaklines, such as Spline or Natural Neighbor.

When interpolating densely spaced data generated by a LIDAR sensor, the advantages and disadvantages of various interpolation techniques may vary from those for photogrammetric data. The most significant difference between LIDAR and photogrammetry is the much higher density of LIDAR pulse returns when compared to masspoints generated by photogrammetry, as illustrated below (Figure 18):

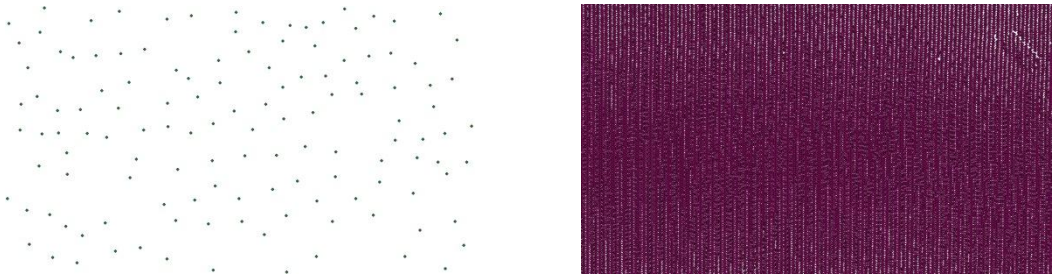

Figure 18 - A comparison between the number of data points typically generated through photogrammetry (left) and the number generated through LIDAR (right). Photogrammetry data points are from the Virginia Base Mapping Program for terrain near Wytheville, Virginia and the Interstate 81/Interstate 77 interchange. The LIDAR data covers the same terrain and was acquired by the Virginia Dept. of Transportation in 2002.

The figure above compares a set of photogrammetrically derived masspoints for Wytheville, Virginia to a LIDAR dataset of the same terrain. The photogrammetric data averages one masspoint every 40 meters, while the LIDAR data averages one laser return every 2.4 meters.

A number of studies comparing interpolation techniques on LIDAR data have found that differences in accuracy diminish with the increased density of the data. Studies by Lloyd and Atkinson (2002) and Chaplot et al. (2006) compared the IDW technique with Kriging and found that, with the high sampling density of LIDAR, neither produced a significantly more accurate

DEM. Similarly, Goncalves (2006) found that all interpolation techniques approach a plateau in their accuracy as point density increases and therefore no technique produces significantly more accurate results for LIDAR data. Liu et al. (2007) conclude that since no interpolation technique is superior for dense datasets, IDW is preferable because it is the easiest and least computationally intense method of interpolation.

In addition to the interpolation technique used to generate the DEM, one must also consider error introduced by varying the resolution of the raster grid. Because interpolation techniques approximate cell values between scattered data points, a higher resolution raster results in more grid cells to approximate. With the interpolation of LIDAR data, however, the increased density of sample points enables rasters to be created with much higher resolution than less dense photogrammetric datasets. As Liu (2008, p. 40-41) points out, "It is inappropriate to generate a high-resolution DEM with very sparse terrain data: any surface so generated is more likely to represent the shape of the specific interpolator used than that of the target terrain because interpolation artifacts will abound . . . On the other hand, generating a low-resolution DEM from high-density terrain data will devalue the accuracy of the original data."

McCullagh (1988) argued that the resolution of a raster DEM should correspond to the density of the sample data to achieve the most realistic terrain model. The density of the data, or the average spacing of each data point, may be determined by taking the square root of the area divided by the number of points in the dataset, or $\sqrt{(A/n)}$, where A is the area covered by the dataset and n is the total number of data points (Liu, 2008). Studies by Behan (2000) and Smith et al. (2003) find that interpolated LIDAR terrain models are, in fact, most accurate when the grid resolution matches the density of the sample points. The interpolation technique used, however, does not appear to affect the optimal grid size, with the exception of the natural

neighbor technique (Smith et al., 2003). Additionally, Goncalves (2006) found that the RMSEs of interpolated LIDAR remained constant with the IDW and Kriging methods regardless of grid size.

Questions arise, however, when generating a raster DEM with a resolution higher than the sample point density. Both Behan (2000) and Liu (2008) warn against increasing the grid resolution beyond the average point spacing due to the possibility of introducing artifacts generated by the interpolation method. Likewise, Smith et al. (2003) found that increasing resolution beyond the average point spacing increased the level of error in LIDAR derived terrain models. Research by Hengl (2006), however, suggests that increasing the resolution to one-half or one-fourth the average point spacing does not introduce statistically significant error and may actually produce a more optimal terrain model.

2.8 The Difference Between Viewshed Predictions and Real-World Visibility

When creating a model of any geographic phenomenon, the model inevitably has certain limitations since it cannot account for all variables and possible scenarios that may exist in the real world. Visibility calculations and viewsheds have similar limitations caused by a host of other factors that impact visibility besides the terrain itself. Trees, buildings and any kind of vegetation or structures may interfere with visibility. The height of the observer or the target object may play a role by affecting the angle of the line-of-sight. Additional factors include weather conditions such as haze or fog, the size and color of the object, as well as the color and texture of any vegetation or terrain behind the target object (Fisher, 1994). Whether an object rises above the skyline, whether it is moving, and the diminished size due to distance from the observer also impact visibility (Fisher, 1996; Bishop, 2002). The accuracy of a viewshed,

therefore, depends not simply on the accuracy of the terrain model upon which it is based, but on how accurately it models the real-world factors that affect visibility.

Some of the earliest research into viewshed models came from Fisher (1991; 1992) who first noted that viewsheds are significantly limited by the "binary" assumption made about visibility. Most viewshed algorithms in commercial GIS software simply determine whether a raster grid cell is "visible" or "not visible" and assign a value of one to those cells deemed visible and zero to those deemed not visible. Fisher (1992) noted that many objects in the real world are often partially visible, making the binary nature of visibility predictions a limitation in their accuracy, especially for coarse resolution raster grids. If a DEM has a resolution of 10 meters, for example, each grid cell covers an area of 100 square meters, large enough that a portion of the grid cell may be visible to an observer while another portion is not. Clearly, the binary assumption of most viewshed algorithms limits their ability to appropriately model real-world visibility.

Fisher (1992; 1994) proposed the creation of *probable viewsheds* (sometimes referred to as "fuzzy" viewsheds) in which the visibility algorithm includes a certain amount of random noise when determining the visibility of a grid cell. Fisher (1991) previously proposed estimating random noise in a DEM by using Monte Carlo methods (generating a random number from a normal distribution with a mean of zero and a standard deviation based on the dataset's RMSE). A probable viewshed, therefore, would calculate visibility based on elevation values that include random error within a certain margin of error (plus or minus the dataset's RMSE). The results of the visibility algorithm to a particular grid cell are tallied over a series of iterations to define the *probability* of that cell being visible. If the algorithm calculates the visibility of a cell twenty times, for example, and each calculation determines that cell to be visible to the sight

of the observer, then the final viewshed will predict that the cell is, in fact, visible. However, if a cell is found to be visible ten times and blocked the other ten times, then the cell is considered to be possibly visible. Thus, by running the visibility algorithm numerous times and determining the proportion of times that a cell is determined to be visible, a viewshed may be created in which each cell is graded according to its likeliness of visibility.

Fisher (1992; 1994; 1995) argues that probable viewsheds provide a more realistic prediction of visibility than simple binary viewsheds. Indeed, users of viewsheds will likely want to know which areas of terrain may be partially visible or have uncertain visibility, as opposed to which areas will have definite visibility or lack of visibility. A military commander seeking to conceal troop movements, for example, may choose to limit movement only to areas that are repeatedly found to be concealed from an enemy's location and avoid those areas where concealment is less certain. Probable viewsheds also benefit landscape architects who seek to determine the likeliness of an unsightly structure being visible from a residential area (Fisher, 1995; 1996).

A set of other questions arise when considering what factors impact the real-world visibility of terrain. If the observer is viewing an object offset from the ground, such as a building, tower or other structure, the height of that object will need to be considered in the visibility algorithm. Also, if the observer is located above the terrain, such as on a building or tower, the observer's height will also need to be included in the calculation. A commonly used example to illustrate this concept is that of a forest fire, where an observer may not be able to see the actual fire, but can see the smoke rising above the fire (Fisher, 1996). Likewise, the tops of buildings, cell phone towers, or other structures may be visible while their bases are concealed behind a ridge in the terrain. Therefore, an accurate viewshed algorithm must consider whether

an object at a particular location rises above the terrain rather than simply considering whether the ground at that location is visible.

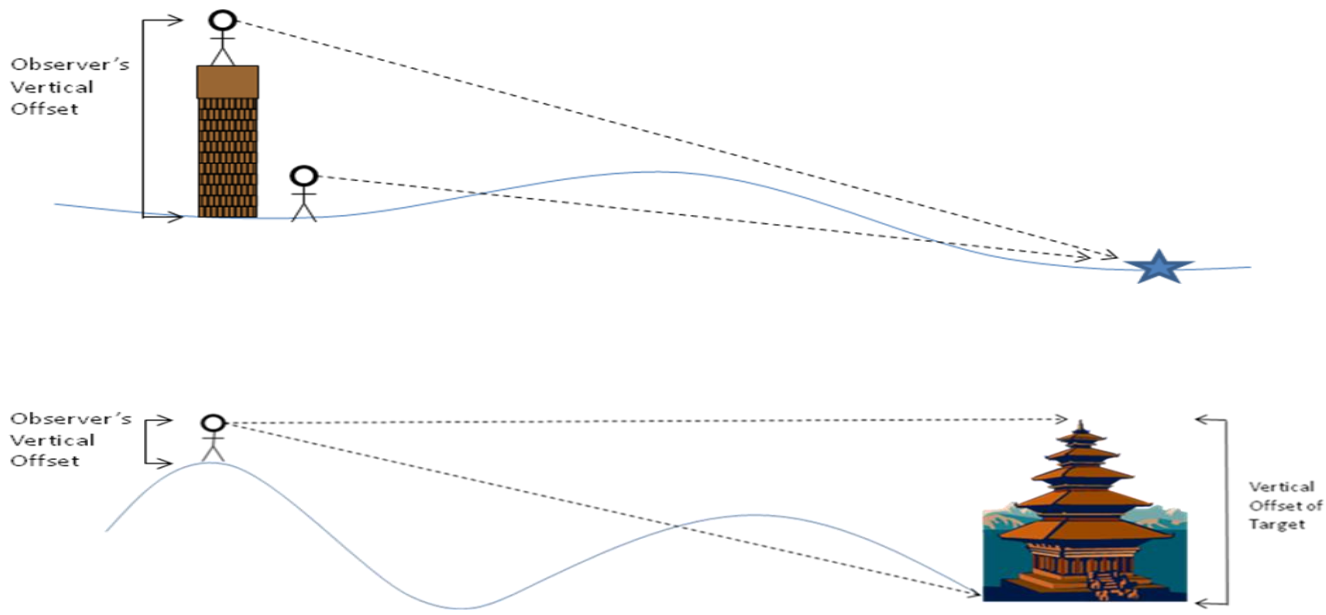

Figure 19 - Offsets for the observer or target often impact visibility. An observer on an elevated position (top) will be able to see targets concealed to an observer on the ground. An elevated target, such as a building (bottom), may rise above the terrain, allowing an observer to see the top but not the base.

In addition to determining the visibility of a particular grid cell, one must also consider whether a person viewing an object at that location will actually notice it. Due to the diminishing apparent size of objects farther away, Fisher (1994) notes that an object may be imperceptible to an observer depending on its distance. This raises the question of what role *perceptibility* may play in addition to visibility when modeling a viewshed. Bishop (2002) describes the concept of a *visibility threshold* when landscape planners seek to minimize the aesthetic impact of a wind turbine farm. Bishop defines several factors affecting the visibility threshold including perceived size of the turbines, their contrast with the surroundings, the effect

of atmospheric scattering, and the fact that wind turbines are often in motion. He concludes that all factors affect the perceptibility of objects by observers, but distance of the object and its corresponding decrease in apparent size have the strongest impact.

2.9 Visibility Through Vegetation

One of the most difficult factors in a viewshed to quantify is the effect of vegetation on visibility. This question is especially pertinent due to the fact that many GIS users generate viewsheds based on bare-earth DEMs which do not include any surface features such as buildings or trees. With the advances made in LIDAR technology, however, generating viewsheds based on LIDAR-derived DSMs is becoming more common. While features such as buildings can block a line-of-sight in the same manner as a hill or ridge, vegetation's impact on the viewshed varies due to the different degrees of visibility *through* vegetation.

Dean (1997) was the first to attempt to measure what he calls the *visual permeability* of vegetation. Just as Fisher (1991; 1992) pointed out that the binary assumption of "visible" vs. "not visible" was a serious shortcoming of viewsheds, Dean notes that viewshed algorithms are similarly limited by making a binary assumption about the permeability of obstructions to the line-of-sight. Viewsheds determine each grid cell along the LOS to be either completely transparent, in which case the LOS is not blocked, or completely opaque, thereby determining the target to be not visible. When the LOS passes through cells that lie on terrain covered by forest, however, visual permeability must be taken into consideration to allow for the possibility that the vegetation may not entirely obstruct the observer's view of the target. Additionally, the amount of obstruction caused by vegetation is cumulative, and therefore the more vegetation between the observer and target, the more the visibility will be obstructed.

Figure 20 - An observer's visibility is cumulatively diminished as the LOS passes through vegetation.

Dean (1997) proposes an algorithm that considers the effect of vegetation by subtracting a certain percentage of visibility for every vegetated grid cell through which the LOS passes. Each LOS, therefore, is initially considered to be 100% visible. With each vegetated grid cell the LOS encounters, a certain amount (called the permeability coefficient) is subtracted from that percentage. If the percentage reaches zero, then the target is deemed to be blocked from view. Otherwise, if an intermediate percentage remains when the LOS reaches the target, then the target is considered to have partially visibility.

Dean (1997) notes, however, that the permeability coefficient varies depending on the type of vegetation; tree species, canopy height, understory growth, and the seasons are all factors that affect the degree of visual permeability. Given a sufficient amount of field data, an index could be created that defines the permeability coefficients for different vegetation types. Llobera (2007), however, proposes an algorithm based on Beer-Lambert's Attenuation Law which is used to model a particle crossing a medium in the study of physics, optics, and meteorology. Llobera (2007) also points out several shortcomings of Dean's proposed algorithm, namely the lack of consideration of the density of vegetation, the angle at which the LOS enters and leaves the vegetation, and the assumption that permeability decreases linearly with greater distance through vegetation when, in fact, permeability should decrease exponentially.

Chapter 3: Methodology

3.1 Study Site and Datasets

The study area chosen for the project is near the town of Wytheville in Wythe County, southwest Virginia. The site was chosen due to the availability of LIDAR data acquired by the Virginia Department of Transportation in 2002. The data was acquired to study the feasibility of relocating the interchange of Interstate 81 and Interstate 77, which lies just to the east of Wytheville. The area lies in the foothills of the Blue Ridge Mountains while Interstate 81 follows an area of low-lying terrain between two ridgelines.

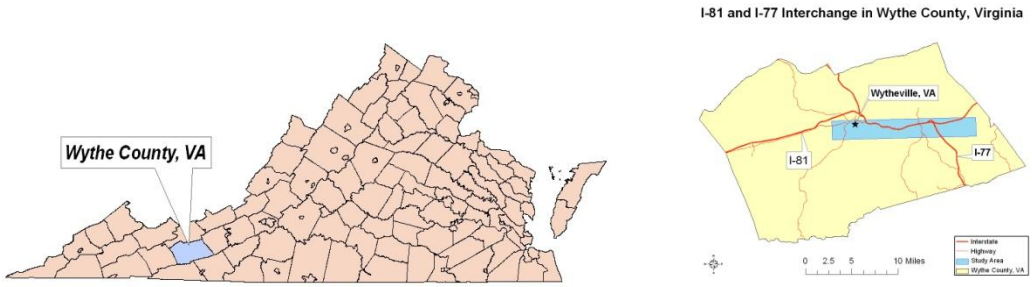

Figure 21 - The study area consists of Wytheville, Virginia and the interchange of Interstate 81 and Interstate 77 in the southwest Blue Ridge Mountains.

The I-81 and I-77 interchange approximately follows Reed Creek which flows through a valley between higher elevation ridgelines that include Sand Mountain and Lick Mountain, just south of Wytheville, along with Ramsey Mountain and Hamilton Knob, on the northeast boundary of the study area. The resulting terrain of the study area, therefore, consists of relatively flat, low-lying land along the Interstates which is mostly farmland, pastures, or small

residential subdivisions. The terrain becomes rougher and increases in elevation along the boundary of the study site near the mountain ridges.

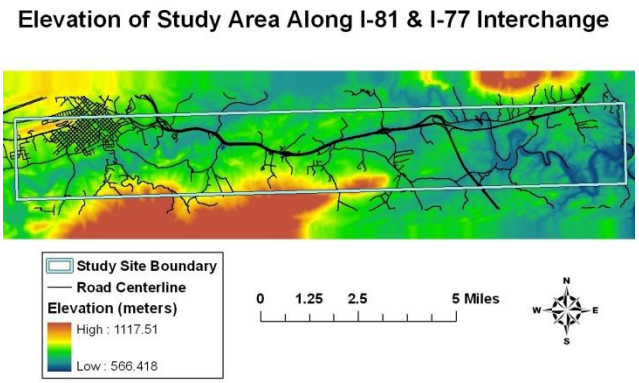

Figure 22 - The elevation of the area near the I-81 & I-77 Interchange overlain with the study area boundary and road centerlines (Elevation data from the Virginia Base Mapping Program; Road Centerline data from the Virginia Dept. of Transportation).

According to personal communication with Raleigh Cook (2009) of the Virginia Department of Transportation, the LIDAR dataset was acquired at a flying height of 4500 feet above sea level at a speed of 200 feet per second. The laser's frequency was 14 Hz with a pulse repetition rate of 10 KHz. The sensor's maximum scanning angle was +/- 16 degrees from the nadir, resulting in a swath width of 2581 feet. The final dataset has an average point spacing of 7.3 feet (2.2 meters), but certain areas have denser point spacing due to overlapping flight lines. The exact Root Mean Square Error of the dataset was unknown.

Certain areas of the terrain, however, suffered from a lack of data points. These include bodies of water such as lakes and streams because water has poor reflectivity of LIDAR pulses. Other areas had missing strips where the sensor had not recorded data along a particular part of a flight line. Some flight lines were also longer than others, leading to some data points extending

beyond the main body of the dataset. These excess points were cut from the main dataset prior to processing into DSMs and areas of missing data were excluded from use when selecting field observation sites. In addition, the LIDAR dataset included only first-returns, reflecting the height of surface features, but obviously limiting the ability to model the vertical structure of forested or vegetation-covered terrain.

The photogrammetric data used in this study was acquired from the Virginia Base Mapping Program (VBMP) for Wythe County, Virginia. The VBMP is managed by the Virginia Geographic Information Network which is a division of the Virginia Information Technologies Agency (VITA). According to the VITA website (2010), the VBMP first acquired high-resolution aerial orthophotos for the entire state of Virginia in 2002 and updated the photos in 2006 – 2007. The orthophotos were used to derive topographic contours spaced at one foot (1:2400 scale) or one-half foot (1:1200 scale), depending on population density and the desires of municipal governments. Digital Terrain Models including both masspoints and breaklines were provided in Microstation DGN format while color aerial photography was available in GeoTIFF and MrSID formats.

Wythe County, Virginia, was one of the municipalities mapped by the VBMP at 1:1200 scale in 2006. The VBMP data for Wythe County included Microstation DGN files with masspoints and breaklines along with aerial orthophotos in MrSID format. The data was divided into 5000' x 5000' tiles and only the tiles that fall inside the boundary of the study area were used. The data's original projection was the Virginia State Plane South coordinate system but was reprojected into UTM Zone 17N to be consistent with the LIDAR dataset's projection.

Figure 23 - Aerial photo (left) and masspoints with breaklines (right) from the Virginia Base Mapping Program near Wytheville, VA.

3.2 Creating Surfaces from the Data

To turn the point datasets from LIDAR and the VBMP into DEMs and DSMs, they needed to be interpolated into surfaces. Because a number of interpolation techniques exist with various advantages and disadvantages, the choice of interpolation technique is a major factor when conducting a study that compares surfaces created from different datasets. The chosen interpolation technique needs to create surfaces that represent the real-world terrain as closely as possible while also minimizing any differences that may be introduced when interpolating two datasets with different characteristics.

Past research has demonstrated that differences in accuracy between interpolation techniques diminish as the point density of the dataset increases, making the choice of interpolation technique for LIDAR largely irrelevant at certain densities (Goncalves, 2006; Liu et al., 2007). In addition, Lloyd and Atkinson (2002) and Chaplot et al. (2006) found no significant differences in accuracy when LIDAR datasets where interpolated with the Inverse Distance

Weighted and Kriging methods. The choice of interpolation technique for this study, therefore, depended on the characteristics of the VBMP data. Past research has also demonstrated that DEMs generated from photogrammetric datasets are significantly more accurate when breaklines are included in the interpolation process (Li, 1992; 1994; Gong et al., 2000). Because the VBMP dataset includes breaklines, an interpolation technique capable of using breaklines would be superior to techniques that do not. In ESRI ArcGIS 9.3, only the IDW and Kriging techniques allow the use of breaklines, making them preferred choices over Spline and Natural Neighbor which only use point data (as of the release of ArcGIS 10, Spline does include a barrier option).

The IDW method was the final choice for interpolating the datasets due to its ability to include the VBMP breakline data in the interpolation process. While Kriging may also use breakline data, a major disadvantage is its computational intensity when compared with IDW. An attempt to interpolate a single 5000' x 5000' tile of VBMP data with breaklines required approximately twelve hours of processing time using the Kriging method with its default settings. IDW, however, required approximately an hour for each tile. Given that the study area covered fifty-seven tiles of VBMP data, a very large amount of time and computing power would be required for the Kriging method. Also, an attempt was made to process a same size tile of the LIDAR data with Kriging using default settings which resulted in a freezing of the ArcGIS software.

In addition to the interpolation technique, the resolution of the raster grids to be used in the analysis needed to be considered. To examine whether the relationship between the resolution of the interpolated surface and the average point spacing of the dataset had an impact on viewshed accuracy, multiple raster grids with varying resolutions were generated from the LIDAR and photogrammetric datasets. The LIDAR dataset averaged one point every 2.2 meters

while the photogrammetric dataset averaged one point every 40 meters. Therefore, raster grids were generated from the LIDAR dataset with 0.5-meter, 1-meter, 2-meter, 5-meter, and 10-meter resolutions while rasters were generated from the photogrammetric data with resolutions of 1-meter, 5-meters, 10-meters, 20-meters, 30-meters and 40-meters. If an interpolated surface is most accurate when its resolution matches the average grid spacing of the point dataset, then the most accurate raster grids should be the 2-meter LIDAR and the 40-meter photogrammetric surfaces. If Hengl's (2006) suggestion that resolutions of one-half or one-fourth the average point spacing are acceptable, then the viewsheds based on the LIDAR surfaces with 1-meter and 0.5-meter resolutions should have comparable accuracy to viewsheds from the 2-meter surface. Likewise, the accuracy of viewsheds based on the photogrammetric DEMs with 20-meter and 10-meter resolutions should be comparable to that of the DEM with 40-meter resolution. The DEMs of 10, 5, and 1-meter resolution may also reveal a decrease in accuracy when the resolution is far higher than the average point spacing. Also, any decrease in viewshed accuracy for surfaces with coarser resolutions than the point dataset should be apparent for the LIDAR surfaces with 5 and 10-meter resolutions.

3.3 Conducting the Field Survey

Before conducting the field survey, the set of observation and target points needed to be chosen. A set of observation points was chosen at random and a set of target points for each observation point was also chosen. Ideally, points would be chosen at random throughout the study area. Complete point randomness, however, was impractical due to the study area including large amounts of private property that is inaccessible without permission of the

landowners. Unfortunately, this limited the selection of observation points to publicly accessible areas such as roads or parks.

The observation points were still chosen randomly from a subset of the region using a GIS layer obtained from the Virginia Department of Transportation showing all road centerlines for Wythe County, Virginia. The layer was clipped to include only road centerlines within the study area and a 10 meter buffer was created around the centerlines. A set of thirty random points was generated within the buffered roads layer to serve as the observation points for the study. Several observation points fell in areas impractical for use, such as rivers or ponds, and new random points had to be generated as replacements. If an observation point fell in the middle of a road, it was moved to the nearest roadside location. The target points were then chosen around each observation point and were often selected from nearby landmarks such as houses, street intersections, water towers, or bridges. Approximately five target points were chosen for each observation point, with the number varying slightly due to the availability of suitable targets for each location. A final total of 30 observation points and 135 target points were used in the study.

The field survey consisted of going to each observation point and determining the real-world visibility of the targets. The visibility of each target was recorded as being either completely visible, mostly visible, partially visible, slightly visible, or completely blocked depending on how much of the target could be seen from the observation point. The reason for the target being blocked from view was also recorded. These included complete obstructions based on topography or man-made structures, as well as partial obstructions due to trees or bushes.

3.4 Determining Visibility from the DEMs and DSMs

To test the visibility of the observation and target points on the DEMs and DSMs, a custom toolbar was programmed for ESRI ArcGIS using Visual Basic .NET and ESRI's ArcObjects. The toolbar allows the user to select target and observation points by either clicking a target cursor on their locations or by entering their easting and northing coordinates into an entry form. The use of the form allows more specificity and enables the user to enter the coordinates for a particular pixel, guaranteeing the visibility test is performed on the same target and observation points each time.

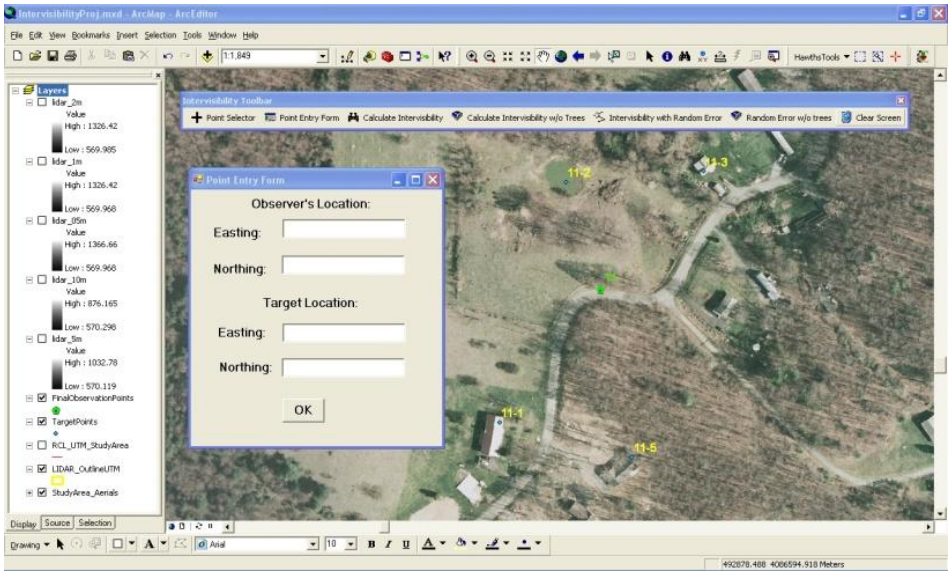

Figure 24 - A custom toolbar was designed for ArcGIS to determine the visibility of target points from observation points. An entry form allows for the input of the easting and northing coordinates in the UTM coordinate system.

The program determines visibility by following a custom programmed algorithm based on the principle of determining whether any terrain rises above the line-of-sight. When the observer's elevation is compared with the target's elevation, and the LOS extends from the

observer to the target, it forms the hypotenuse of a triangle between the two locations, allowing for comparison of the LOS angle to the angles between the observer and intermediate grid cells. If any terrain in between the observer and target has a greater elevation than the elevation of the LOS, the target point is determined to be blocked. If no intermediate terrain rises above the elevation of the LOS, the target point is considered to be visible.

The algorithm first determines several features relating to the observer and target locations: the number of rows in the raster grid between the observer and target, the number of columns between the observer and target, and the change in elevation of the LOS. The difference in rows and columns is the same as the change in the X, Y grid coordinates of the pixels separating the observer and target.

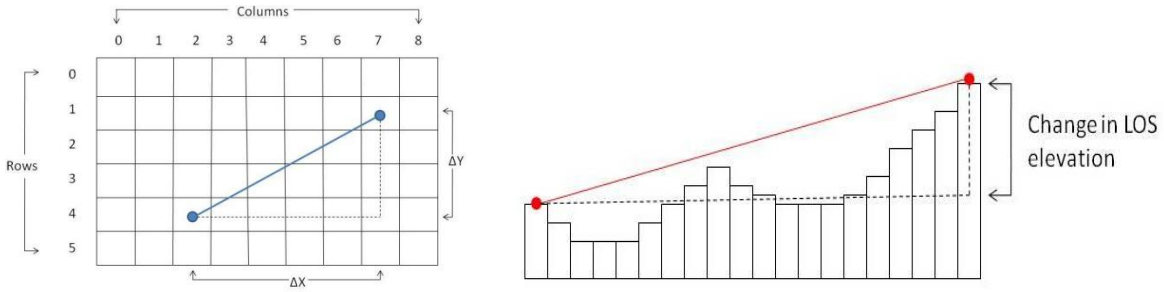

Figure 25 - The algorithm determines the difference between the row and column locations of the observer and target (ΔX and ΔY, plan view at left) as well as the change in elevation for the line-of-sight (cross-section view at right).

The difference in elevation between observer and target is divided by the longer dimension between them (either ΔX or ΔY) to determine the rate at which the LOS changes elevation per pixel of the longer dimension. The ratio of the dimensions (ΔX / ΔY) is also determined to allow for the algorithm to calculate the points at which the LOS crosses

boundaries between rows and columns. The elevation at the points of a boundary crossing can then be interpolated from the neighboring pixels.

The elevation values obtained from the raster grid are loaded into an array, representing the elevation of the terrain underneath the LOS. Each elevation value in the array is then compared to the corresponding elevation of the LOS at the same location. Any time the terrain's elevation exceeds that of the line-of-sight's elevation, the target pixel is reported to be blocked from the observer's sight.

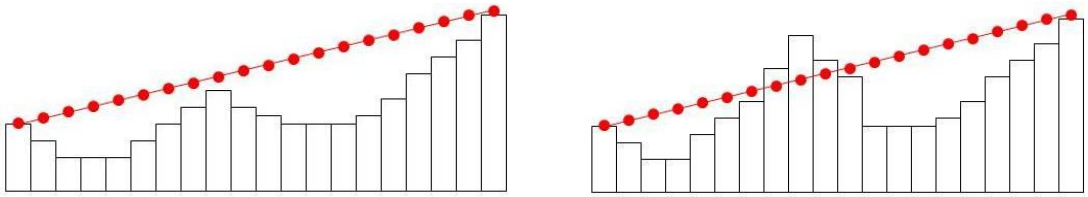

Figure 26 - An array holds the elevation values for the LOS (represented by red dots). If the elevation of the LOS is constantly higher than the underlying terrain, the target is visible (left). If the elevation of an underlying pixel exceeds the elevation of the corresponding point in the LOS, the target is blocked (right).

The program also includes the option of adding an offset to the observer and target for cases where either is elevated above the ground. Offset values are added to the elevation values for the pixels in which the observer and target are located. If the target is a building or some other structure or landmark, its maximum height needs to be included as an offset when calculating visibility with a bare-earth DEM. When using a LIDAR DSM, however, no offset is required since the LIDAR surface already includes structures and buildings. An offset for the observer, however, may be included to account for the height of the observer's eyes.

3.5 Adding Random Error to Determine *Probable Visibility*

As noted by Fisher (1991; 1992; 1994) a shortcoming of viewshed algorithms is the labeling of a target as being either visible or not visible and can be improved by determining the *probable visibility* of a target. This involves running the algorithm through multiple iterations and adding an amount of random error to the elevation values beneath the line-of-sight. Because the viewshed is based on DEMs or DSMs which are generated from datasets including a certain amount of error, the random error is based on the Root Mean Square Error of the dataset. The error is normally distributed with zero as the mean and one-half of the RMSE as the standard deviation.

To determine the probable visibility, the program follows Fisher's (1991) suggestion to use the Box-Muller transform (Box & Muller, 1958) to generate a normally distributed series of random numbers based on the elevation dataset's RMSE (A Visual Basic implementation of the Box-Muller transform can be found at Xtreme Visual Basic Talk, 2006). Since 95% of values under a normal distribution are within +/- two standard deviations of the mean, the Box-Muller transform, using a standard deviation equal to one-half the dataset's RMSE, generates random numbers with 95% being within the RMSE from the mean at zero. The program adds a random error value to each number in the array of elevation values under the line-of-sight, thereby creating a new elevation array that includes random error.

The visibility algorithm is performed one thousand times using a new line-of-sight elevation array with random values generated each time. The results for the iterations of the algorithm are summarized and the percentage of times the target is deemed visible and not visible is reported. The probable visibility of the target, therefore, is the percentage of trials in which the target was found to be visible. If the target is visible 100% of the time or blocked

from view 100% of the time, the program has definitively found the target to be either visible or blocked. If the target is visible only part of the time, however, such as a case in which the target is reported visible only 40% to 60% of the time, then the user should be cautious about drawing a conclusion about the target's visibility in the real world. The results of the visibility calculations with random error will be compared to the results without random error to determine whether it yields a more realistic prediction when compared with the results of the field survey.

The RMSE for the LIDAR surfaces was set at 15 centimeters since this is within the generally expected RMSE for most LIDAR data (Fowler, 2007). An exact RMSE for the original LIDAR dataset was not available. The RMSE for the DEM surfaces was determined using a set of Ground Control Points (GCPs) surveyed by the National Geodetic Survey. The GCPs were located at various points throughout the study area with many lying near major roads such as the interstates. The elevation value at each control point was compared with the value of each interpolated surface at the same location to calculation the RMSE for that surface. The following RMSEs were determined for the different resolution DEMs:

Surface	RMSE (meters)
DEM 1m	4.07
DEM 5m	4.15
DEM 10m	4.15
DEM 20m	4.00
DEM 30m	4.06
DEM 40m	4.03

3.6 Eliminating Trees from the LIDAR DSMs for Comparison with Bare-Earth DEMs

The most significant difference between viewsheds based on LIDAR-derived surfaces and those from bare-earth photogrammetry is the inclusion of surface features such as trees and buildings in the LIDAR. This inevitably leads to viewsheds based on LIDAR having an advantage in accuracy for targets obstructed by surface features. Moreover, the issue of vegetation between the observer and target raises the possibility that a target may be deemed not visible by the program, yet in reality, may be visible *through* the vegetation, an option not supported in traditional GIS applications.

To appropriately compare the advantage of LIDAR DSMs to bare-earth DEMs, a second series of visibility analyses was performed on the LIDAR surfaces with all vegetated areas in the LIDAR surface delineated. Any time the algorithm determined that the line-of-sight was blocked at a particular location, the program determined whether that point fell within an area populated by trees or other vegetation. If so, that point was not considered to be a point of blockage, and the algorithm continued along the line-of-sight, discarding all points of blockage that fell within vegetated areas. If the LOS was blocked only in areas with vegetation, then the target was labeled as being visible to the observer. Therefore, targets were determined to be blocked from sight only if the location of the blockage fell in an area not covered by vegetation. By excluding the vegetated areas, the visibility analyses on the LIDAR surfaces may be more accurately compared to the bare-earth DEMs based on the photogrammetric data.

Figure 27 – By delineating the areas of terrain covered by trees and vegetation, a mask can be used to remove visibility blockages caused by vegetation from the analysis. A shapefile outlining the vegetated areas is overlain on the VBMP aerial photo (left) and the LIDAR-derived raster (right). The green triangle marks the observer's location and the blue dots mark the target locations.

3.8 Estimating Visibility Based on the Amount of Vegetation Along the LOS

The final analysis for the project is to determine whether a relationship exists between the length of the LOS that lies in the vegetated terrain and the degree of visibility determined in the field survey. Using the target points which at least have some degree of visibility in the field, the distance that the LOS passes through the vegetated layer was compared to the degrees of real-world visibility in the field using logistic regression. One would expect that the greater the distance of the LOS that passes through vegetated terrain would cause the target to be less visible. A regression was performed between the two variables to determine the extent to which this occurs.

The logistic regression between the two variables approximates Dean's (1997) proposal to estimate visibility of a target by subtracting a certain percentage of visibility for each vegetated pixel through which the LOS passes. Since the field survey ranked targets according to their level of visibility, each target obscured by vegetation was ranked on a scale of zero to three according to whether it was mostly visible, partially visible, or slightly visible or

completely blocked. Targets mostly visible were assigned a ranking of three while targets completely blocked were ranked as zero. Slightly visible and partially visible targets were ranked as one and two respectively. For those targets with partial visibility due to vegetation, the ordinal regression between the visibility rankings and the distance that the LOS passes through vegetation should produce a means to estimate the visibility of a target based on the amount of vegetation between it and the observer.

Chapter 4: Analysis and Results

4.1 Degrees of Visibility in the Field Survey

The real-world visibility of the targets depended upon the *degree* to which they were visible from the observer's location as some targets were partially obscured by terrain, vegetation or man-made structures. As noted in the methodology, all targets in the field survey were labeled as being completely visible, mostly visible, partially visible, slightly visible, or completely blocked depending on how much of the target could be seen. A target with no obstructions would be considered completely visible while a target with minor obstructions, such as a house with small trees in front, would be considered mostly visible. When approximately 50% of the target could be seen, it was labeled partially visible, while a target that was mostly obscured, but at least some part could be identified, was considered slightly visible.

Figure 28 - Varying degrees of visibility depend on the amount of obstruction between the target and observer. For mostly visible targets such as the house at left, the entire target can be seen but is slightly obscured by light vegetation. Partially visible targets such as the pond (center) have significantly more obstruction but are still visible. Slightly visible targets such as the house at right can be identified, but the majority of the structure is obscured by vegetation.

The determination of the degree of visibility in the field survey introduces an obvious source of subjectivity to the analysis. In addition, when comparing the results of the visibility algorithm to the field survey, the accuracy of the algorithm's predictions depend upon which

degree of visibility is considered necessary for the target to be considered truly visible in the real world. As noted in the methodology, the original visibility algorithm simply labels the targets as being either visible or not visible, with no ability to determine any levels of partial visibility. The accuracy of the algorithm when compared with the field study, therefore, depends upon whether a target must meet a minimum threshold of visibility to be defined as being truly visible. If one considers only targets that are fully visible in the field to be defined as visible, then the algorithm produces correct predictions when it finds only those targets to be visible and finds all other targets to be blocked. If one defines visibility as being any target that is at least partially or slightly visible, however, then the algorithm yields a correct prediction when it predicts those targets as being visible and only considers completely blocked targets to be not visible.

As an example, the visibility algorithm may predict a building will be visible but the field survey may find that building to be mostly blocked from sight with the exception of a small portion of the roof. Should the algorithm's prediction be considered accurate? On the other hand, if the visibility algorithm should predict that building to be blocked, should the prediction be considered correct since the majority of that building would not be visible to the observer? If one considers any degree of real-world visibility as qualifying a target to be visible, then the visibility algorithm must determine the building to be visible in order to produce a correct prediction. However, if one considers a target to be truly visible only if 50% or more of it is visible, then the visibility algorithm would yield a correct prediction only if it determines the building to be blocked.

Because the accuracy of the visibility algorithm varies depending on the real-world visibility of the targets, the results of the algorithm had to be classified accordingly when being compared to the results of the field survey. The proportions of algorithm's predictions

considered correct or incorrect depend on what levels of real-world visibility are defined as being visible. Therefore, four separate comparisons of the algorithm's predictions with the field results were produced, each defining different degrees of real-world visibility as meeting the threshold for being considered visible. The first comparison considered targets with any level of real-world visibility to be visible, meaning the algorithm yields a correct prediction whenever it determines a target is visible and any portion of that target was visible in the field. On the other hand, if the algorithm predicts a target to be not visible, the prediction is correct only when the field survey finds that target to be completely blocked from view. The second comparison considered partially visible, mostly visible, and completely visible targets to be considered truly visible while targets that were only slightly visible were considered to be blocked from view. Likewise, the third comparison considered only mostly visible and completely visible targets to be visible while targets that were partially visible or slightly visible were considered blocked from view. The final comparison considered only targets that were fully visible in the field as being visible when compared with the predictions of the visibility algorithm. The four different comparisons are illustrated in Table 1:

Comparison 1			Comparison 2		
	Considered Visible	Considered Not Visible		Considered Visible	Considered Not Visible
Completely Visible	**X**		Completely Visible	**X**	
Mostly Visible	**X**		Mostly Visible	**X**	
Partially Visible	**X**		Partially Visible	**X**	
Slightly Visible	**X**		Slightly Visible		**X**
Completely Blocked		**X**	Completely Blocked		**X**
Comparison 3			**Comparison 4**		
	Considered Visible	Considered Not Visible		Considered Visible	Considered Not Visible
Completely Visible	**X**		Completely Visible	**X**	
Mostly Visible	**X**		Mostly Visible		**X**
Partially Visible		**X**	Partially Visible		**X**
Slightly Visible		**X**	Slightly Visible		**X**
Completely Blocked		**X**	Completely Blocked		**X**

Table 1 - Four sets of comparisons were made between the field survey results and the visibility algorithm's predictions depending on which visibility rankings were defined as being visible in the field survey.

When the visibility algorithm's predictions are compared with the results of the field survey, another issue that arises is the difference between false positives and false negatives. False positives occur when the algorithm predicts the target to be visible but, in the field, it is not visible. A false negative occurs when the program predicts the target to be blocked from view, yet in reality the target is visible. The proportion of incorrect predictions that are false positives and false negatives also varies depending on which degrees of real-world visibility are defined as being visible. In the comparisons above, the more field verified visibility levels that are considered visible, the more we would expect false positives. Similarly, in comparisons that are more stringent, we expect more false negatives.

4.2 The Effect of Resolution on Visibility Accuracy

While the accuracy of the visibility analysis varied based on resolution, not all datasets showed a particular trend towards higher accuracy at certain resolutions. Past research indicates that interpolated surfaces are most accurate when the cell size matches the average point spacing of the dataset, which can be determined by taking the square root of the area covered divided by the number of data points (the formula $\sqrt{(A/n)}$). For the LIDAR, this formula yielded an optimal resolution of approximately two meters while the optimal resolution of the VBMP photogrammetric data was 40 meters. If the optimal resolution for the interpolated surfaces yields the best result for visibility analysis, then the visibility algorithm should have the highest accuracy when applied to the 2 meter LIDAR DSM and 40 meter VBMP DEM.

For most of the analyses with the DSMs and DEMs, there does not appear to be a significant overall trend for the different resolutions. For the LIDAR surfaces of 0.5, 1, and 2 meters, the accuracy levels for the visibility algorithm are within 2-3 percentage points. When only fully visible real-world targets are considered, the algorithm yielded an accuracy of 71.0% for all three surfaces. With other degrees of visibility, the visibility algorithm is slightly more accurate for the 0.5 and 1-meter DSMs than the 2-meter DSM. When only mostly visible and fully visible targets are considered, for example, the algorithm had an accuracy of 68.4% for the 2-meter DSM and an accuracy of 71.0% for the 0.5 and 1-meter DSMs. For the other degrees of real-world visibility, the algorithm is also slightly more accurate on the 0.5 and 1-meter DSMs than on the 2-meter surface. When all levels of field visibility are considered, the algorithm's accuracy was 60.6% on the 2-meter surface while on the 0.5 and 1-meter surfaces, it had accuracies of 61.9% and 62.6% respectively. When only partially, mostly, and fully visible

targets were considered, the algorithm had an accuracy of 65.8% on the 2-meter DSM while it had accuracies of 67.1% and 68.4% for the 0.5 and 1-meter DSMs, respectively.

Decreasing the resolution of the LIDAR surfaces, however, did tend to decrease the accuracy of the visibility algorithm. The algorithm consistently had the lowest accuracy levels for the 10-meter LIDAR DSM of all the LIDAR surfaces. The difference in the algorithm's accuracy between the 10-meter surface and the 0.5, 1, and 2-meter surfaces increased as more degrees of field visibility were considered. When only fully visible targets were considered, the algorithm had an accuracy of only 63.2% for the 10-meter DSM compared with 71.0% for the 0.5, 1, and 2-meter DSMs. The gap in the visibility algorithm's accuracy rates between the 10-meter grid and higher resolution grids decreased as more degrees of real-world visibility were considered, but when applied to the 10-meter DSM, the algorithm consistently had accuracy rates of at least 3-4 percentage points below the highest accuracy for all DSMs and all degrees of real-world visibility.

Degrees of field visibility considered:		***All visibility levels***		***Partially, mostly, and fully visibile***		***Mostly and fully visible***		***Only fully visible***	
	Surface Resolution	% Correct	% Incorrect	% Correct	% Incorrect	% Correct	% Incorrect	% Correct	% Incorrect
	LIDAR 0.5m	61.9	38.1	67.1	32.9	71.0	29.0	71.0	29.0
LIDAR DSMs	LIDAR 1m	62.6	37.4	68.4	31.6	71.0	29.0	71.0	29.0
	LIDAR 2m	60.6	39.4	65.8	34.2	68.4	31.6	71.0	29.0
	LIDAR 5m	62.6	37.4	69.0	31.0	67.7	32.3	67.7	32.3
	LIDAR 10m	59.4	40.6	64.5	35.5	64.5	35.5	63.2	36.8

Table 2 - The visibility algorithm on the LIDAR DSMs had comparable accuracies at 0.5, 1, and 2-meter resolutions. The algorithm on the 10-meter DSM, however, had consistently lower accuracy.

Overall, the visibility algorithm tended to have the highest accuracy for the 1-meter LIDAR DSMs. Because the difference in accuracies tended to be small, however, an element of random chance may play a role when the visibility algorithm has a higher accuracy on the 1-meter DSM than the 0.5 or 2-meter DSMs. However, based on this study's results, it appears that a slightly higher resolution DSM does not decrease the accuracy of visibility analysis based on LIDAR data. Using a LIDAR DSM with one-half or one-fourth the optimal resolution (as defined by $\sqrt{(A/n)}$), therefore, should provide a statistically similar visibility accuracy as an optimal resolution DSM. It may be in the user's interest, however, to use the optimal resolution DSM in order to reduce the amount of time and computing power required to perform a visibility analysis.

The visibility algorithm's accuracy on the DSMs with less-than-optimal resolutions, however, does decrease. The algorithm had the lowest accuracy rate on the 10-meter LIDAR DSM for all resolutions across all degrees of real-world visibility. Because the algorithm was consistently less accurate for 10-meter DSM than the 0.5, 1, or 2-meter DSMs, conducting visibility analysis on surfaces with lower resolution than the LIDAR's average point spacing would not be recommended.

The impact of resolution for the photogrammetric DEMs, however, was less clear. Overall, the visibility algorithm had much lower accuracy for the DEMs than the LIDAR DSMs, with accuracy rates often being no more than 50% when all degrees of visibility were considered. When only fully visible targets were considered, the accuracy dropped to around 40%. The DEMs also do not show a clear pattern of improved accuracy at the optimal resolution of 40 meters, with the algorithm's accuracy rates on the 1 and 5-meter DEMs either tying or exceeding the accuracy on the 40-meter surface for all degrees of visibility. The algorithm had the best

accuracy overall for the 5-meter DEM with either the highest accuracy or tying for the highest accuracy for all degrees of visibility, except when mostly or fully visible targets were considered. The algorithm on the 30-meter DEM was the least accurate overall as it had either the lowest accuracy or was tied for the lowest accuracy for all degrees of visibility considered.

Degrees of field visibility considered:		***All visibility levels***		***Partially, mostly, and fully visibile***		***Mostly and fully visible***		***Only fully visible***	
	Surface Resolution	*% Correct*	*% Incorrect*	*% Correct*	*% Incorrect*	*% Correct*	*% Incorrect*	*% Correct*	*% Incorrect*
	DEM 1m	53.5	46.5	52.9	47.1	48.4	51.6	41.9	58.1
	DEM 5m	54.8	45.2	52.9	47.1	48.4	51.6	41.9	58.1
DEMs	DEM 10m	50.3	49.7	52.9	47.1	49.7	50.3	41.3	58.7
	DEM 20m	51.6	48.4	51.6	48.4	45.2	54.8	40.0	60.0
	DEM 30m	50.3	49.7	51.6	48.4	45.2	54.8	40.0	60.0
	DEM 40m	52.9	47.1	52.9	47.1	46.5	53.5	41.3	58.7

Table 3 - The visibility algorithm's accuracies for different resolution DEMs did not show as consistent a pattern as the LIDAR DSMs. The algorithm had the highest overall accuracy on the 5-meter surface and the lowest overall accuracy on the 30-meter surface. There appears to be no discernable relation between the algorithm's accuracy and the optimal resolution.

Since the visibility algorithm's accuracy on the photogrammetric DEMs does not appear to relate to the optimal resolution, creating a viewshed from a DEM with a higher-than-optimal resolution should not decrease the accuracy. However, since the algorithm on the DEMs had an accuracy of close to 50%, regardless of resolution, the user faces the risk of the results being largely due to random chance. Clearly, the algorithm achieved a higher accuracy on the LIDAR DSMs than the photogrammetric DEMs making LIDAR's use preferable to photogrammetric data when conducting a visibility analysis. If a DEM is to be used to predict visibility, however,

it is clearly most accurate when any degree of target visibility is considered, as requiring more stringent degrees of visibility makes reduces the visibility accuracy to as low as 40.0%.

4.3 LIDAR DSMs and Photogrammetric DEMs

For all further comparisons of the LIDAR DSMs and the photogrammetric DEMs, the most accurate surface resolutions were used. The visibility algorithm had the most accurate results overall on the 1-meter surface for the LIDAR and the 5-meter DEM. These two surfaces were compared when analyzing the algorithm's rates of false positives and false negatives, the addition of random error to the elevation values, and the impact of vegetation on the targets' visibility.

As would be expected with their incorporation of surface features, the visibility algorithm had a greater accuracy for LIDAR-generated DSMs than the bare-earth DEMs derived from photogrammetry. The algorithm on the 1-meter LIDAR surface had an accuracy of 71.0% when only targets that were fully visible in the field were considered visible. This was the highest accuracy of all the LIDAR surfaces of all resolutions. The algorithm on the 5-meter bare-earth DEM had the highest accuracy for all the DEMs when all degrees of visibility were considered, at 54.8%. It is clear, therefore, that the increased accuracy of the LIDAR yields significantly more accurate visibility predictions than a photogrammetric DEM.

In addition to the visibility algorithm's differing rates of accuracy between the DSMs and the DEMs, a significant difference occurred in the proportions of false positives and false negatives. The algorithm was more likely to produce false negatives on the LIDAR DSMs, where a target is predicted to be blocked from sight, but is visible in the field. For the 1-meter LIDAR surface, 70.7% of the algorithm's errors were false negatives when all degrees of field

visibility are considered. The proportion of false negatives declines as the degrees of visibility become more stringent, to a low of 17.7% when only fully visible targets are considered. Clearly, for the LIDAR DSMs, the visibility algorithm errs more often by predicting a target to be blocked when, in reality, the target is at least slightly or partially visible. Many of these errors were found to be due to vegetation, where a target would be partially obscured, yet still identifiable, by trees or bushes.

Higher proportions of false positives occurred consistently on the 5-meter bare-earth DEM. When only fully visible targets were considered, the algorithm's incorrect predictions for the 5-meter DEM were 95.6% false positives with 4.4% false negatives. The proportion of false positives declines as more degrees of field visibility are considered. When any degree of visibility is considered visible, the algorithm had a false positive rate of 70.0% on 5-meter DEM and a false negative rate of 30.0%. Because the DEMs were bare-earth models, higher false positive rates are to be expected since the algorithm is more likely to predict a target to be visible when, in reality, the target is blocked by surface features.

Degrees of field visibility considered:	***All visibility levels***		***Partially, mostly, and fully visibile***		***Mostly and fully visible***		***Only fully visible***	
Surface Resolution	*% False Positive*	*% False Negative*	*% False Positive*	*% False Negative*	*% False Positive*	*% False Negative*	*% False Positive*	*% False Negative*
LIDAR 1m	29.3	70.7	38.8	61.2	64.4	35.6	84.4	15.6
DEM 5m	70.0	30.0	77.8	22.2	90.0	10.0	95.6	4.4

Table 4 - False positive and false negative rates for the algorithm on the 1-meter LIDAR and 5-meter DEM with different visibility levels considered. The algorithm on the LIDAR surface has consistently higher rates of false negatives while the algorithm on bare-earth DEM has higher rates of false positives.

The results show that the algorithm always has higher rates of false negatives on the LIDAR DSMs and higher rates of false positives on the DEMs. This holds true for all

resolutions and for all degrees of real-world visibility considered. This primarily results from the LIDAR DSMs including surface features such as vegetation and buildings, while the photogrammetric data for the DEMs does not. The algorithm, therefore, is more likely to predict false positives on the DEMs because of their bare-earth representation of the terrain when, in reality, the targets are often obstructed by the surface features.

The algorithm on the LIDAR DSMs, however, produced a higher proportion of false negatives due to obstructions caused by the surface features. When conducting the field survey, it became clear that the majority of false negatives on the LIDAR surfaces are due to vegetation. Many targets that were partially visible in the field were predicted to be blocked for the LIDAR DSMs, arising from the visibility algorithm's inability to consider the visual permeability of vegetation. Whenever the visibility algorithm determined a pixel between the observer and target rose above the line-of-sight's elevation, that target was determined to be blocked from sight, regardless of whether that pixel may have been caused by vegetation that, in the field, does not prevent the observer from identifying the target. Areas of vegetation on the DSM, therefore, are treated by the visibility algorithm as being entirely opaque when, in fact, the observer can often see targets through the vegetation. The following photos illustrate such a situation, in which targets were found to be visible through the vegetation in the field, yet the visibility algorithm determined them to be blocked from sight on the LIDAR DSMs:

Figure 29 – Some partially obscured targets were determined to be blocked with the LIDAR DSMs although they are still identifiable to the observer in the field. These include a house (left), a barn (center), and a pond (right). These same targets were determined to be visible on the bare-earth DEMs.

4.4 Results with Random Error

In addition to predicting the visibility of the targets based on the interpolated LIDAR and photogrammetric surfaces, an analysis was conducted in which an element of random error was added to the elevation of each pixel along the line-of-sight between the observer and the target. An algorithm implementing the Box-Muller transform was used to generate normally distributed random error with a standard deviation of one-half the dataset's Root Mean Square Error, making 95% of the random error values within +/- the dataset's RMSE. An error value was added to each elevation value between the observer and the target and visibility was determined using the new elevation values. The program repeated this process one thousand times and then calculated the percentage of times the target was deemed visible or blocked.

When comparing the results of the visibility calculations with random error to those without, the target was judged to be visible if the LOS was blocked less than 50% of the time. If the LOS was blocked more than 50% of the time, the target was considered to be blocked from the observer's view. If the target was visible more than 50% of the time and that same target was determined to be visible without the addition of random error, then the predictions with and without the random error were considered to be in agreement. If the prediction with random

error differed from the prediction without random error, then it was recorded as being a disagreement.

The LIDAR surfaces were found to clearly be less susceptible to having a different visibility prediction due to the addition of random error. For the 1-meter LIDAR DSM, the visibility algorithm's predictions with random error agreed with the original predictions 96.8% of the time. The algorithm on the 5-meter DEM, however, had a lower rate of agreement with only 85.8% of the predictions with random error being the same as those without. The most likely explanation for this result is the LIDAR's higher accuracy, given that a 15 centimeter RMSE was used for the LIDAR surfaces, as opposed to a RMSE of 4.15 meters for the 5-meter DEM.

	Agree	Disagree	% Agree	% Disagree
LIDAR 1m	150	5	96.8	3.2
DEM 5m	133	22	85.8	14.2

Table 5 - The number of times the addition of random error generated a prediction that agreed or disagreed with the original prediction that did not include random error.

4.5 Results for LIDAR DSMs with Vegetated Terrain Masked

Because the vegetation clearly resulted in a high proportion of false negatives for the LIDAR DSMs, a second analysis was conducted with a mask over the pixels that lie in areas with significant vegetation cover. A shapefile was created for each field sample whose line-of-sight traversed vegetated terrain between the observation point and its corresponding target points. With the visibility algorithm modified to discard any blockages to the line-of-sight that fell within the vegetation polygons, the visibility of the targets could be predicted without the

excessive rates of false negatives caused by the vegetation. This essentially treats all vegetated terrain as if it were entirely transparent, whereas in the original LIDAR surfaces, any vegetation that blocks the line-of-sight is assumed to be opaque.

With the vegetated terrain masked out, the number of false negatives declined significantly when compared to the original LIDAR DSMs. The proportion of false positives, however, increased sharply to levels similar to those of the bare-earth DEMs. When all degrees of real-world visibility were considered visible, the predictions for the 1-meter LIDAR DSM with vegetated areas masked had a false positive rate of 64.3% while the original 1-meter LIDAR DSM had a false positive rate of 29.3%. The predictions for the two surfaces had false negative rates of 35.7% and 70.7%, respectively. The visibility algorithm's false positive rate for the 1-meter LIDAR with the vegetation mask was lower than that of the 5-meter bare-earth DEM, on which it had a rate of 70.0%. Likewise, the predictions for the DEM had a lower false negative rate than for the LIDAR with the vegetation mask, at 30.0%. Table 6 compares the algorithm results for the original 1-meter LIDAR DSM, the 1-meter LIDAR DSM with the vegetation mask, and the 5-meter bare-earth DEM:

	Match	Mismatch	% Match	% Mismatch	False Positive	False Negative	% False Positive	% False Negative
LIDAR 1m	97	58	62.6	37.4	17	41	29.3	70.7
LIDAR 1m w/ veg. mask	99	56	63.9	36.1	36	20	64.3	35.7
DEM 5m	85	70	54.8	45.2	49	21	70	30

Table 6 - Masking out vegetated terrain from the LIDAR DSMs reduces the number of false negatives, but increases the number of false positives. However, the rate of false positives for the LIDAR remains less than that of the bare-earth DEM. The results consider all degrees of field visibility to be visible.

The overall accuracy of the algorithm on the LIDAR DSMs with vegetation masked was similar to that of the original LIDAR DSMs when all degrees of field visibility were considered visible. However, the predictions for the LIDAR with the vegetation mask were still more accurate than for the bare-earth DEMs. The algorithm had an accuracy of 63.9% for the 1-meter DSM with the vegetation mask compared to 54.8% for the 5-meter DEM. The algorithm's accuracy on the LIDAR with the vegetation mask decreased, however, as the degrees of visibility became more stringent. When only fully visible targets are considered visible, the algorithm had an accuracy of only 49.7% on the 1-meter DSM with the vegetation mask compared to the original 1-meter DSM on which it had an accuracy of 71.0%. Clearly, the predictions for the LIDAR with the vegetation mask follow a similar pattern as those for the bare-earth DEMs in that they are most accurate when all degrees of field visibility are considered and decrease in accuracy as the degrees of field visibility are more restricted.

	Match	Mismatch	% Match	% Mismatch	False Positive	False Negative	% False Positive	% False Negative
LIDAR 1m	110	45	71.0	29.0	38	7	84.4	15.6
LIDAR 1m w/ veg. mask	77	78	49.7	50.3	74	4	94.9	5.1
DEM 5m	65	90	41.9	58.1	86	4	95.6	4.4

Table 7 - Results when only fully visible targets are considered visible.

The algorithm's different accuracy rates between the original LIDAR and the LIDAR with the vegetation mask are due to the lower proportion of false negatives that occur when the vegetation is masked. The predictions for the original LIDAR had much higher rates of false

negatives caused by vegetation while those for the LIDAR with the vegetation mask had higher rates of false positives. It is clear, therefore, that when any degree of real-world visibility is considered, a LIDAR DSM with vegetation removed may yield a more accurate estimate by reducing the number of false negatives caused by permeable vegetation. When the criteria for real-world visibility becomes more restricted, as in cases where a target must be fully visible, the accuracy of the algorithm on LIDAR without vegetation decreases and the use of original LIDAR DSMs is superior.

When comparing the algorithm on the LIDAR DSMs with a vegetation mask to the bare-earth DEMs, however, the LIDAR is clearly superior. For all degrees of real-world visibility, the predictions for the LIDAR DSMs without vegetation had higher accuracy rates than those for the bare-earth DEMs, usually between 8 – 12 percentage points higher. Therefore, if the user desires higher overall accuracy, LIDAR is clearly preferable, either with or without vegetation, over bare-earth DEMs generated from photogrammetric data.

4.6 Accuracy for Targets with No Surface Features Along the Line-of-Sight

As a comparison between the visibility accuracy of the LIDAR DSMs and the bare-earth DEMs, all target points which did not have any buildings or vegetation along their lines-of-sight were selected and their accuracy was calculated. Since the lines-of-sight for these targets are free of surface features, they most closely resemble the results one would expect from an analysis performed on a bare-earth model derived from the LIDAR dataset. A total of 37 targets of the set of 135 had lines-of-sight that were not obstructed by surface features.

The accuracy of the visibility algorithm for the selected targets tended to increase as the degrees of visibility were restricted, following a similar pattern as the bare-earth DEMs. The

highest accuracy was 67.6% on the 0.5-meter LIDAR DSM when only fully visible targets were considered visible or when mostly and fully visible targets were considered visible. The lowest accuracy was 45.9% on the 10-meter DSM with all degrees of visibility considered visible.

The accuracy of the visibility algorithm for targets with no surface features along their lines-of-sight also tended to decrease with coarser resolution. When the visibility algorithm was applied to the DSMs for these targets, the 0.5-meter DSM had the highest accuracy, or was tied for the highest accuracy, for all degrees of visibility. The algorithm on the 10-meter DSM consistently had the lowest accuracy for these targets.

Degrees of field visibility considered:		***All visibility levels***		***Partially, mostly, and fully visibile***		***Mostly and fully visible***		***Only fully visible***	
	Surface Resolution	% Correct	% Incorrect	% Correct	% Incorrect	% Correct	% Incorrect	% Correct	% Incorrect
Targets w/ no surface features on LOS	LIDAR 0.5m	54.1	45.9	59.6	40.4	67.6	32.4	67.6	32.4
	LIDAR 1m	54.1	45.9	56.8	43.2	64.9	35.1	62.2	37.8
	LIDAR 2m	51.4	48.6	54.1	45.9	62.2	37.8	59.5	40.5
	LIDAR 5m	51.4	48.6	54.1	45.9	62.2	37.8	59.5	40.5
	LIDAR 10m	45.9	54.1	48.6	51.4	56.8	43.2	54.1	45.9

Table 8 - Accuracy for the visibility algorithm on the LIDAR DSMs for a subset of 37 of the targets which had no surface features along their lines-of-sight.

4.7 Resolution and LIDAR DSMs with Vegetation Mask

For the LIDAR DSMs with the vegetated terrain masked, the lower resolution surfaces increase in prediction accuracy over the higher resolution surfaces. For all degrees of real-world visibility, the predictions for the 10-meter DSM had the highest accuracy with the vegetation mask. Additionally, the predictions for the 10-meter DSM also had a lower number of false positives than the higher resolution grids. Because the original LIDAR DSMs suffered from a

high rate of false negatives due to permeable vegetation, removing the vegetated terrain results in lower rates of false negatives, leaving the rate of false positives as the greater influence on the total accuracy. Decreasing the number of false positives, therefore, results in a higher overall accuracy. For the LIDAR DSMs with vegetation masked out, decreasing the resolution has the opposite effect when compared with the original DSMs and actually appears to increase the accuracy. A possible explanation for this result may be the lesser number of pixels between the observer and target for coarser resolution surfaces.

Degrees of field visibility considered:		***All visibility levels***		***Partially, mostly, and fully visibile***		***Mostly and fully visible***		***Only fully visible***	
	Surface Resolution	% Correct	% Incorrect	% Correct	% Incorrect	% Correct	% Incorrect	% Correct	% Incorrect
LIDAR with vegetation mask	LIDAR 0.5m	62.6	37.4	61.3	38.7	56.1	43.9	49.7	50.3
	LIDAR 1m	63.9	36.1	62.6	37.4	56.1	43.9	49.7	50.3
	LIDAR 2m	62.6	37.4	61.3	38.7	57.4	42.6	51.0	49.0
	LIDAR 5m	61.3	38.7	62.6	37.4	57.4	42.6	50.3	49.7
	LIDAR 10m	65.2	34.8	65.2	34.8	60.0	40.0	56.1	43.9

Table 9 - The visibility algorithm on the LIDAR DSMs with the vegetation mask had the highest accuracy rate for the 10-meter surface.

4.8 Accuracy for Bare-Earth DEMS

One clear trend for the DEMs is the tendency for visibility accuracy to decrease as the degrees of real-world visibility become more restricted. When all degrees of real-world visibility are considered, the DEMs have prediction accuracies between 50.3% - 54.8%. When only fully visible targets are considered, however, the accuracy falls to 40.0% - 41.9%. The decrease in accuracy is largely caused by the tendency of bare-earth DEMs to have a higher number of false positives. When the degrees of real-world visibility are more restricted, the number of false

positives increases, thereby decreasing the overall accuracy rates for the DEMs. When using bare-earth surfaces for visibility analysis, therefore, the results will most closely match the real world if targets are considered visible anytime any portion of the target is visible, no matter how slight.

4.9 Estimation of Visibility Through Vegetation Using Logistic Regression

The final goal of this project was to determine a means of estimating visibility through vegetation. A logistic regression was performed between the degree of visibility of the targets obscured by vegetation and the distance their lines-of-sight travel through that vegetation. This idea is based on a proposal by Dean (1997), who suggested decrementing the visibility by a certain percentage for every pixel for which the LOS passed through vegetation. Using the vegetation layers created for the analysis of the LIDAR surfaces with a vegetation mask, the length of the LOS that crosses terrain covered in vegetation could be calculated. This distance was equal to or less than the total distance between the observer and target, since not all terrain between observer and target included vegetation. Some targets had a long LOS, but only a small length of the LOS passed through vegetation, while others had much greater LOS distances through vegetation, depending on the proportion of the LOS that passed through vegetated terrain. None of the observation points were in vegetation.

Of the 155 original targets, seventy were partially or fully obscured by vegetation. The degree of visibility observed in the field was ranked accordingly, with 0 = not visible, 1 = slightly visible, 2 = partially visible, and 3 = mostly visible. Of the seventy target points, thirty-five were randomly selected as a test data set. A logistic regression was performed between the

visibility ranking and the distance the LOS passed through vegetation. The results of the regression were then used to predict the visibility of the remaining thirty-five targets. A contingency table was then produced to compare the predicted visibilities of the second set of targets with their field results.

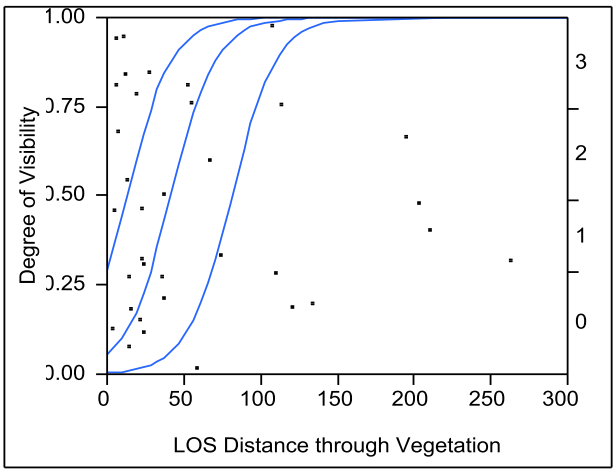

Figure 30 - Logistic Regression plot between the distance the LOS passes through vegetation and the degree of visibility of the targets in the field survey.

As one would suspect, a greater amount of vegetation between the observer and target makes the target less likely to be visible. On the set of randomly selected thirty-five targets, the logistic regression found several breakpoints between the visibility rankings, when mostly visible becomes partially visible, partially visible becomes slightly visible, and slightly visible becomes not visible. A weighted average of the LOS distances was calculated to determine the exact break point when the model predicts a change in the visibility categories. All targets with an LOS distance through vegetation of more than 74.49 meters were found to be blocked from sight, while a distance of 36.87 to 74.48 meters resulted in the target being slightly visible. A

distance of 17.73 to 36.86 meters resulted in the target being partially visible while distances of less than 17.72 meters were found to be mostly visible.

Distance Of LOS through Vegetation (meters)	Predicted Result by Logistic Regression
74.49 and above	0 - Not Visible
36.87 - 74.48	1 - Slightly Visible
17.73 - 36.86	2 - Partially Visible
0 - 17.72	3 - Mostly Visible

Table 10 - Break points between the visibility rankings predicted by the logistic regression.

The break points determined by the logistic regression on the first set of thirty-five points were then tested on the second set of thirty-five target points. A contingency analysis was performed using the model's visibility predictions for the second set and the field results. An error matrix was produced and the statistical significance of the relationship between the field visibility and the distance the LOS passed through vegetation was determined.

The results of the contingency analysis on the second set of thirty-five targets showed that the logistic regression did find a significant relationship between the distance the LOS passed through vegetation and the field visibility of the target. As the distance the LOS passes through vegetation, the degree of visibility of the target declines. The analysis resulted in a Kappa statistic of 0.559497 with a probability of being greater than Chi Squared of less than 0.0001. The error table produced by the contingency analysis is shown below:

		Degree of Visibility (in Field)				
	Count Total % Col % Row %	0	1	2	3	
Model Prediction	0	13 37.14 82.86 100.00	0 0.00 0.00 0.00	0 0.00 0.00 0.00	0 0.00 0.00 0.00	13 37.14
	1	1 2.86 7.14 33.33	1 2.86 50.00 33.33	0 0.00 0.00 0.00	1 2.86 14.29 33.33	3 8.57
	2	0 0.00 0.00 0.00	0 0.00 0.00 0.00	5 14.29 41.67 83.33	1 2.86 14.29 16.67	6 17.14
	3	0 0.00 0.00 0.00	1 2.86 50.00 7.69	7 20.00 58.33 53.85	5.00 14.29 71.43 38.46	13 37.14
		14 40.00	2 5.71	12 34.29	7 20.00	35

Table 11 - The error table resulting from a contingency analysis between the logistic regression's visibility predictions and the field results for the second set of thirty-five targets. The analysis produced a Kappa statistic of 0.559497 with a Prob>ChiSq of less than 0.0001.

Based on the analysis there clearly exists a significant correlation between the distance the observer-to-target LOS passed through vegetation and the visibility of that target, at least for the vegetation in this region at this time of year. It is therefore reasonable to predict that a target will be less visible with a greater linear distance of vegetation between it and the observer. While visibility through vegetation is also affected by factors such as the vegetation's density,

leaf size, and seasonal changes, it is clear that the distance the LOS passes through vegetation is an important component in any attempt to model visibility and one that is easily determined by the algorithm using only the GIS data layers provided to produce the surface and locate the vegetation. With further research to quantify other factors, it may be possible to develop a model that yields reasonable predictions as to the visual permeability of vegetation.

Chapter 5: Conclusions

The results of this study clearly show that, overall, LIDAR surfaces surpass those derived from photogrammetric data in visibility accuracy. When applied to the LIDAR surfaces, the visibility algorithm generated accuracy rates as high as 71.0% while the lowest accuracy rate was 59.4%. For the photogrammetric DEMs, however, the visibility algorithm rarely had accuracy rates higher than 50%, with the highest being only 54.8%. Clearly, the most accurate DEM surface is inferior to the least accurate LIDAR surface at predicting visibility. Masking out the vegetated terrain on the LIDAR surfaces reduced the algorithm's accuracy when compared with the unmasked LIDAR, but was still more accurate than when applied to the bare-earth DEMs.

The visibility algorithm's greater accuracy on the LIDAR surfaces may be attributed to the greater accuracy of the elevation measurements, given that LIDAR often has an RMSE of about 15 centimeters, while the photogrammetric surfaces had RMSEs of approximately four meters. The greater point density of the LIDAR dataset also contributes to creating a more accurate interpolated surface because the higher density reduces the influence of interpolation technique as pointed out by Goncalves (2006). Even when the vegetated terrain is masked out of the LIDAR surface, the lower RMSE and higher point density would explain the greater accuracy of the visibility predictions over those for the photogrammetric DEMs.

Whether a user should conduct visibility analyses with a DSM with all earth surface features or a bare-earth model depends on the user's needs. The use of the visibility algorithm on the LIDAR DSMs that include vegetation clearly produce more false negatives due to the algorithm's inability to consider the permeability of the vegetation. If the user desires to conceal

targets from an observer's sight, then using a bare-earth model would be superior to a DSM since it would be less likely to err due to vegetation. Locating the target on terrain deemed not visible with a bare-earth model would guarantee that concealment is due to topography and not vegetation which may not fully conceal the target from the observer. If the user desires that a target be visible to an observer, however, using a LIDAR DSM would be the best choice since it would be more accurate in selecting locations that are not obstructed by terrain, vegetation, or man-made structures.

If LIDAR data is available for terrain being analyzed, its use should be preferred to that of photogrammetric data. If the user desires to generate viewsheds that take into consideration vegetation and man-made structures, then LIDAR is definitely more suited to the task than bare-earth photogrammetric data. If the user desires a bare-earth model for analysis, however, LIDAR data with surface features removed would likely be superior to bare-earth photogrammetric data, due to LIDAR's lower RMSE and the higher density of data points.

The resolution of the interpolated surface does not appear to have a significant impact on the accuracy of the visibility predictions except when the resolution is coarser than the optimal resolution. When the resolution of the LIDAR DSMs is increased to one-half or one-fourth the optimal resolution the accuracy remains comparably similar. The accuracy rates decrease only when the resolution is less than the optimal resolution, such as for the 5-meter and 10-meter LIDAR surfaces. When the vegetated terrain is masked out on the LIDAR surfaces, however, the algorithm's accuracy does tend to decrease with finer-than-optimal resolutions and is highest for the coarser 10-meter surface. For the DEMs, the accuracy rates are comparable for all resolutions, but the accuracy is usually close to 50%, the same result one would get through random chance.

When random error is added to the elevation values between the observer and target, the prediction results for the LIDAR surfaces are largely similar. Only a small number of targets had differing predictions when random error was added as opposed to the original LIDAR. Therefore, the visibility algorithm's predictions on the LIDAR surfaces appear to be less affected by the addition of random error. The predictions for the DEMs, however, tended to have a higher proportion that differed when random error was added. The difference between the DEMs and LIDAR DSMs is likely due to the fact that the LIDAR has a lower RMSE of only 15 centimeters while the DEMs all had RMSEs of approximately four meters. A greater RMSE likely contributes to a greater tendency to alter the result when error is added. In addition, the proportion of predictions that differed when adding random error increased as the resolution of the DEMs increased, possibly due to the greater number of pixels that lie in between the observer and target for finer resolution DEMs.

In addition, the logistic regression performed between the distance the LOS passes through vegetation and the visibility of the targets in the field survey shows that one can estimate visibility of a target based on the amount of vegetation between the target and observer. While a number of different factors play a role in determining visibility through vegetation, the length of the LOS through vegetated terrain clearly does correlate to the likelihood of the target being visible. Also, the logistic regression showed that there is a certain threshold (approximately seventy meters) at which the target is likely to be entirely blocked from view. Just as Dean (1997) predicted that the effect of vegetation on visibility would be cumulative, the logistic regression suggests that the probability of a target being visible will experience a plateau effect, where any targets with a length of LOS through vegetation greater than a particular threshold may be reliably predicted to be blocked from view.

5.1 Ideas for Further Research

This study shows the difficulty of modeling visibility in GIS with the numerous factors that affect visibility in the real world. The most significant obstacle to accurate visibility analysis with the LIDAR DSMs is the inability of the algorithm to consider the visual permeability of the vegetation that may partially, but not completely, obstruct the observer's view of the target. Various proposals have been made by Dean (1997) and Llobera (2007) to approach the issue of visual permeability, which both rely on the ability to measure and quantify the degree of the vegetation's permeability. While this study showed that the extent to which vegetation obscures a target can be related to the distance that the LOS passes through vegetation, visual permeability also involves a number of other factors. The density and type of vegetation, its type of leaves (i.e. whether it is deciduous or coniferous trees), and whether the vegetation cover varies with the seasons all play a role in the degree to which it may obscure the visibility of a target. Measuring these factors would be further complicated in cases with mixed types of vegetation, such as a forest with significant understory growth. Using 360-degree ground-based LIDAR may be able to determine the degree of permeability for vegetation in a particular area. Understanding the amount of visual permeability for vegetated terrain, however, will be vital to improving the accuracy of GIS visibility predictions.

In addition to the possibility of using ground based LIDAR, aerial LIDAR with multiple point returns may also prove helpful in modeling visibility through vegetation. The LIDAR dataset for this study included only first returns, making the surfaces generated from the data reflect only the top of the vegetation cover. LIDAR data with multiple returns, however, can be used to determine different strata in multi-layered vegetation, such as the difference between a forest canopy and the height of understory vegetation. The ability to model the different

vegetation layers would allow the user to develop differing estimates on the visual permeability of the vegetation based on which layer through which the LOS passes. If the LOS were to pass through a forested area with a thick canopy but sparse understory growth, for example, the likely visibility of the target would depend on whether the LOS is at the level of the canopy or the understory.

Another factor which would likely increase the reliability of visibility studies would be to consider the entire footprint of a target when considering that target's visibility. Many of the targets in this study that were buildings or structures were often partially visible due to one section being visible but other parts being blocked from sight. Therefore, when a target is a building or structure, calculating the visibility of all the pixels in the target's footprint would result in a more accurate prediction of whether that target is visible to the observer.

When predicting visibility based on Monte Carlo methods, the ideal viewshed would divide the pixels into categories such as quartiles or percentiles. Rating pixels on whether they were visible 0-10% of the time, 10-20% of the time, and so on, would create a viewshed in which each pixel is graded on the likelihood of it being visible to the observer. The selection of location, then, would depend on the user's needs. If the user seeks to conceal an object from the location of an observer, it would best be located on terrain that has a 0% rate of visibility from the observer's location. If the user desires that a target be visible to an observer, however, the ideal location would be on terrain that has a 100% rate of visibility to the observer's location.

Overall, the difficulty of modeling visual permeability through vegetation illustrates the need for visibility algorithms that can model semi-transparent surfaces. The use of LIDAR datasets requires a new algorithm that considers a degree of transparency to the surface features, especially vegetation that may block the line-of-sight. Since the target can often be seen through

vegetation, terrain with trees or bushes needs to be considered partially transparent and not entirely opaque. With multiple-return LIDAR datasets, an algorithm could be devised to consider factors such as vegetation density and vertical structure when the line-of-sight passes through a set of trees or bushes. Ideally, an algorithm could identify the likelihood that a target would be visible through vegetation in a similar manner as when Monte Carlo methods are used to add random error to the elevation values along the line-of-sight.

Appendix A

Results Depending on Degree of Visibility

Any degree of field visibility considered visible

		Match	Mismatch	% Match	% Mismatch	False Positive (Predicted Visible)	False Negative (Predicted Blocked)	% False Positive	% False Negative
LIDAR DSMs	LIDAR 0.5m	96	59	61.9	38.1	17	42	28.8	71.2
	LIDAR 1m	97	58	62.6	37.4	17	41	29.3	70.7
	LIDAR 2m	94	61	60.6	39.4	20	41	32.8	67.2
	LIDAR 5m	97	58	62.6	37.4	21	37	36.2	63.8
	LIDAR 10m	92	63	59.4	40.6	25	38	39.7	60.3
LIDAR with Vegetation mask	LIDAR 0.5m	97	58	62.6	37.4	36	22	62.1	37.9
	LIDAR 1m	99	56	63.9	36.1	36	20	64.3	35.7
	LIDAR 2m	97	58	62.6	37.4	36	22	62.1	37.9
	LIDAR 5m	95	60	61.3	38.7	37	23	61.7	38.3
	LIDAR 10m	101	54	65.2	34.8	29	25	53.7	46.3
DEMs	DEM 1m	83	72	53.5	46.5	49	23	68.1	31.9
	DEM 5m	85	70	54.8	45.2	49	21	70	30
	DEM 10m	78	77	50.3	49.7	52	25	67.5	32.5
	DEM 20m	80	75	51.6	48.4	52	23	69.3	30.7
	DEM 30m	78	77	50.3	49.7	54	23	70.1	29.9
	DEM 40m	82	73	52.9	47.1	51	22	69.9	30.1

Partially Visible, Mostly Visible and Fully Visible Targets Considered Visible

		Match	Mismatch	% Match	% Mismatch	False Positive (Predicted Visible)	False Negative (Predicted Blocked)	% False Positive	% False Negative
LIDAR DSMs	LIDAR 0.5m	104	51	67.1	32.9	19	32	37.3	62.7
	LIDAR 1m	106	49	68.4	31.6	19	30	38.8	61.2
	LIDAR 2m	102	53	65.8	34.2	22	31	41.5	58.5
	LIDAR 5m	107	48	69.0	31.0	23	25	47.9	52.1
	LIDAR 10m	100	55	64.5	35.5	28	27	49.1	50.9
LIDAR with vegetation mask	LIDAR 0.5m	95	60	61.3	38.7	43	17	71.7	28.3
	LIDAR 1m	97	58	62.6	37.4	43	15	74.1	25.9
	LIDAR 2m	95	60	61.3	38.7	43	17	71.7	28.3
	LIDAR 5m	97	58	62.6	37.4	42	16	72.4	27.6
	LIDAR 10m	101	54	65.2	34.8	35	19	64.8	35.2
DEMs	DEM 1m	83	72	52.9	47.1	56	16	77.8	22.2
	DEM 5m	83	72	52.9	47.1	56	16	77.8	22.2
	DEM 10m	83	72	52.9	47.2	55	17	76.4	23.6
	DEM 20m	80	75	51.6	48.4	59	16	78.7	21.3
	DEM 30m	80	75	51.6	48.4	59	16	78.7	21.3
	DEM 40m	82	73	52.9	47.1	57	16	78.1	21.9

Mostly Visible and Fully Visible Considered Visible

						False Positive	False Negative		
		Match	Mismatch	% Match	% Mismatch	(Predicted Visible)	(Predicted Blocked)	% False Positive	% False Negative
LIDAR DSMs	LIDAR 0.5m	110	45	71.0	29.0	28	17	62.2	37.8
	LIDAR 1m	110	45	71.0	29.0	29	16	64.4	35.6
	LIDAR 2m	106	49	68.4	31.6	32	17	65.3	34.7
	LIDAR 5m	105	50	67.7	32.3	36	14	72	28
	LIDAR 10m	100	55	64.5	35.5	40	15	72.7	27.3
LIDAR with vegetation mask	LIDAR 0.5m	87	68	56.1	43.9	59	9	86.8	13.2
	LIDAR 1m	87	68	56.1	43.9	60	8	88.2	11.8
	LIDAR 2m	89	66	57.4	42.6	58	8	87.9	12.1
	LIDAR 5m	89	66	57.4	42.6	58	8	87.9	12.1
	LIDAR 10m	93	62	60.0	40.0	51	11	82.3	17.7
DEMs	DEM 1m	75	80	48.4	51.6	72	8	90	10
	DEM 5m	75	80	48.4	51.6	72	8	90	10
	DEM 10m	77	78	49.7	50.3	70	8	89.7	10.3
	DEM 20m	70	85	45.2	54.8	76	9	89.4	10.6
	DEM 30m	70	85	45.2	54.8	76	9	89.4	10.6
	DEM 40m	72	83	46.5	53.5	74	9	89.2	10.8

Only Fully Visible Targets Considered Visible

		Match	Mismatch	% Match	% Mismatch	False Positive (Predicted Visible)	False Negative (Predicted Blocked)	% False Positive	% False Negative
LIDAR DSMs	LIDAR 0.5m	110	45	71.0	29.0	37	8	82.2	17.7
	LIDAR 1m	110	45	71.0	29.0	38	7	84.4	15.6
	LIDAR 2m	110	45	71.0	29.0	39	6	86.7	13.3
	LIDAR 5m	105	50	67.7	32.3	45	5	90	10
	LIDAR 10m	98	57	63.2	36.8	50	7	87.7	12.3
LIDAR with vegetation mask	LIDAR 0.5m	77	78	49.7	50.3	73	5	93.6	6.4
	LIDAR 1m	77	78	49.7	50.3	74	4	94.9	5.1
	LIDAR 2m	79	76	51.0	49.0	72	4	94.7	5.3
	LIDAR 5m	78	77	50.3	49.7	72	5	93.5	6.5
	LIDAR 10m	87	73	56.1	43.9	63	10	86.3	13.7
DEMs	DEM 1m	65	90	41.9	58.1	86	4	95.6	4.4
	DEM 5m	65	90	41.9	58.1	86	4	95.6	4.4
	DEM 10m	64	91	41.3	58.7	84	7	92.3	7.7
	DEM 20m	62	93	40.0	60.0	89	4	95.7	4.3
	DEM 30m	62	93	40.0	60.0	89	4	95.7	4.3
	DEM 40m	64	91	41.3	58.7	87	4	95.6	4.4

Appendix B

Rate of Agreement Between Original Visibility Predictions and Predictions with Added Random Error

		Agree	Disagree	% Agree	% Disagree
LIDAR DSMs	LIDAR 0.5m	149	6	96.1	3.9
	LIDAR 1m	150	5	96.8	3.2
	LIDAR 2m	148	7	95.5	4.5
	LIDAR 5m	146	9	94.2	5.8
	LIDAR 10m	149	6	96.1	3.9
LIDAR with vegetation mask	LIDAR 0.5m	143	12	92.3	7.7
	LIDAR 1m	142	13	91.6	8.4
	LIDAR 2m	144	11	92.9	7.1
	LIDAR 5m	140	15	90.3	9.7
	LIDAR 10m	140	15	90.3	9.7
DEMs	DEM 1m	91	64	58.7	41.3
	DEM 5m	133	22	85.8	14.2
	DEM 10m	144	11	92.9	7.1
	DEM 20m	144	11	92.9	7.1
	DEM 30m	148	7	95.5	4.5
	DEM 40m	149	6	96.1	3.9

Appendix C

Results of Logistic Regression Between the Length of the LOS that Passes Through Vegetation and the Degree of Visibility Determined in the Field Survey

Ordinal Logistic Fit for Degree of Visibility
Logistic Plot

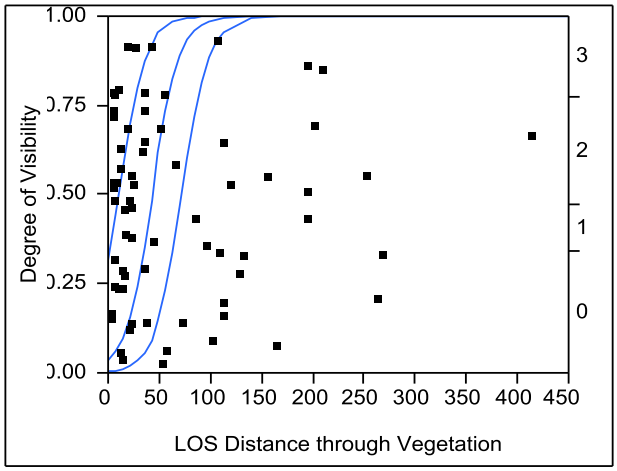

Whole Model Test

Model	-LogLikelihood	DF	ChiSquare	Prob>ChiSq
Difference	42.506898	1	85.0138	<.0001*
Full	50.543065			
Reduced	93.049963			

RSquare (U)	0.4568
Observations (or Sum Wgts)	70
Converged by Objective	

Lack Of Fit

Source	DF	-LogLikelihood	ChiSquare
Lack Of Fit	197	48.633522	97.26704
Saturated	198	1.909543	**Prob>ChiSq**
Fitted	1	50.543065	1.0000

Parameter Estimates

Term		Estimate	Std Error	ChiSquare	Prob>ChiSq
Intercept[0]		-5.5303486	1.0180093	29.51	<.0001*
Intercept[1]		-3.303482	0.6238766	28.04	<.0001*
Intercept[2]		-0.7768936	0.4275607	3.30	0.0692
LOS Distance through Vegetation		0.07679077	0.0147586	27.07	<.0001*

Effect Likelihood Ratio Tests

Source	Nparm	DF	L-R ChiSquare	Prob>ChiSq	
LOS Distance through Vegetation	1	1	85.0137966	<.0001*	

Comparison of field survey results and the Most Likely Degree of Visibility as predicted by the logistic regression where:

0 = Not Visible
1 = Slightly Visible
2 = Partially Visible
3 = Mostly Visible

Degree of Visibility in Field Survey	Distance LOS Passes Through Vegetation	Most Likely Degree of Visibility Predicted by Logistic Regression
0	415.0	0
0	269.5	0
0	262.7	0
0	253.6	0
0	210.4	0
0	203.4	0
0	196.1	0
0	195.9	0
0	195.3	0
0	165.0	0
0	155.9	0
0	133.5	0
0	128.8	0
0	120.2	0

0	113.7	0
0	112.7	0
0	112.1	0
0	110.2	0
1	107.0	0
0	101.4	0
0	96.1	0
0	85.5	0
0	74.2	0
1	66.9	1
0	58.2	1
2	55.2	1
0	53.5	1
2	52.5	1
1	45.0	1
3	42.7	2
1	36.8	2
2	36.7	2
2	36.5	2
2	36.2	2
1	35.3	2
2	34.8	2
3	27.7	2
2	24.3	2
1	23.6	2
2	23.4	2
2	23.0	2
2	22.9	2
2	22.3	2
1	21.9	2
3	19.9	2
3	19.5	2
2	17.6	2
2	15.6	2
2	15.5	2
2	14.7	2
1	14.1	2
2	14.0	2
3	13.4	2

3	11.9	3
1	11.8	3
2	11.4	3
3	10.6	3
3	9.1	3
2	7.7	3
3	7.6	3
2	7.2	3
3	7.2	3
3	7.2	3
3	5.5	3
3	5.5	3
3	5.4	3
3	4.9	3
3	4.5	3
2	4.1	3
2	3.3	3

References

Amidror, I. 2002. Scattered data interpolation methods for electronic imaging systems: a survey. *Journal of Electronic Imaging.* 11(2):157-176.

Bater, C.W. and Coops, N.C. 2009. Evaluating error associated with lidar-derived DEM interpolation. *Computers & Geosciences.* 35:289-300.

Behan, A. 2000. On the matching accuracy of rasterized scanning laser altimeter data. *IAPRS, Vol. XXXIII, Amsterdam.*

Bishop, I.D. 2002. Determination of thresholds of visual impact: the case of wind turbines. *Environmental Planning B: Planning and Design.* 29:707-718.

Bishop, I.D. and Gimblett H.R. 2000. Management of recreational areas: GIS, autonomous agents, and virtual reality. *Environmental Planning B: Planning and Design.* 27:423-435.

Box, G.E.P. and Muller, M.E. 1958. A note on the generation of random normal deviates. *The Annals of Mathematical Statistics.* 29(2):610-611.

Bresenham, J.E. 1965. Algorithm for computer control of a digital plotter. *IBM Systems Journal.* 4(1):25-30.

Cash, J.M. 2003. *Using Light Detection and Ranging (LIDAR) Imagery to model radio wave propagation.* Blacksburg, Va: University Libraries, Virginia Polytechnic Institute and State University. http://scholar.lib.vt.edu/theses/available/etd-04032003-104145.

Chaplot, V., Darboux, F., Bourennane, H., Leguedois, S., Silvera N. and Phachomphon, K. 2006. Accuracy of interpolation techniques for the derivation of digital elevation models in relation to landform types and data density. *Geomorphology.* 77:126-141.

Childs, C. 2004. Interpolating surfaces in ArcGIS Spatial Analyst. *ArcUser*. July-September: 32-35.

Ciciarelli, J.A. 2002. The geology of the battle of Monte Cassino, 1944. *Fields of Battle: Terrain in Military History.* Eds. Doyle, P. and Bennet, M.R. Boston: Kluwer Academic Publishers. 325-344.

Cook, Raleigh. 2009. Personal communication via e-mail. Virginia Department of Transportation.

Dean, D.J. 1997. Improving the accuracy of forest viewsheds using triangulated networks and the visual permeability method. *Canadian Journal of Forest Research.* 27(7):969-977.

Dodd, H.M. 2001. *The validity of using a geographic information system's viewshed function as a predictor for the reception of line-of-sight radio waves.* Blacksburg, Va: University Libraries, Virginia Polytechnic Institute and State University. http://scholar.lib.vt.edu/theses/available/etd-09202001-155212.

Doyle, P., Bennet, M.R., Macleod, R. and Mackay, L. 2002. Terrain and the Messines Ridge, Belgium, 1914-1918. *Fields of Battle: Terrain in Military History.* Eds. Doyle, P. and Bennet, M.R. Boston: Kluwer Academic Publishers. 205-224.

Ehlen, J. and Abrahart, R.J. 2002. Effective use of terrain during the American Civil War: The Battle of Fredericksburg, December 1862. *Fields of Battle: Terrain in Military History.* Eds. Doyle, P. and Bennet, M.R. Boston: Kluwer Academic Publishers. 63-98.

Ehlen, J. and Abrahart, R.J. 2004. Terrain and its affect on the use of artillery in the American Civil War. *Studies in Military Geography and Geology.* Eds. Caldwell, D.R., Ehlen, J., and Harmon, R.S. Boston: Kluwer Academic Publishers. 155-172.

ESRI ArcGIS Resource Center. 2010a. Using viewshed and observer points for visibility analysis. http://help.arcgis.com/en/arcgisdesktop/10.0/help/index.html#/Performing_visibility_analysis_with_Viewshed_and_Observer_Points/009z000000v8000000/.

ESRI ArcGIS Resource Center. 2010b. Understanding interpolation analysis. http://webhelp.esri.com/arcgisdesktop/9.3/index.cfm?TopicName=Understanding_inteinterpola_analysis.

Fisher, P.F. 1991. First experiments in viewshed uncertainty: the accuracy of the viewshed area. *Photogrammetric Engineering & Remote Sensing.* 57(10):1321-1327.

Fisher, P.F. 1992. First experiments in viewshed uncertainty: simulating fuzzy viewsheds. *Photogrammetric Engineering & Remote Sensing.* 58(3):345-352.

Fisher, P.F. 1994. Probability and fuzzy models of the viewshed operation. *Innovations in GIS 1.* Ed. Worboys, M. London: Taylor & Francis. 161-175.

Fisher, P.F. 1995. An exploration of probable viewsheds in landscape planning. *Environment and Planning B: Planning and Design.* 22:527-546.

Fisher, P.F. 1996. Extending the applicability of viewsheds in landscape planning. *Photogrammetric Engineering & Remote Sensing.* 62(11):1297-1302.

Fisher, P.F. 1998. Improved modeling of elevation error with geostatistics. *GeoInformatica.* 2(3):215-233.

Fowler, R. 2007. Topographic Lidar. *Digital Elevation Model Technologies and Applications: The DEM Users Manual.* Ed. Maune, D.F. Bethesda, MD: The American Society for Photogrammetry and Remote Sensing. 207-236.

Garcia-Quijano, M.J., Jensen, J.R., Hodgson, M.E., Hadley, B.C., Gladden, J.B., and Lapine, L.A. 2008. Significance of altitude and posting density on LIDAR-derived elevation accuracy on hazardous waste sites. *Photogrammetric Engineering & Remote Sensing.* 74(9):1137-1146.

Goncalves, G. 2006. Analysis of interpolation errors in urban digital surface models created from Lidar data. *7th International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences.* Eds. Caetano, M. and Painho, M. 160-168.

Gong, J., Li, Z., Zhu, Q., Sui, H. and Zhou Y. 2000. Effects of various factors on the accuracy of DEMs: an intensive experimental investigation. *Photogrammetric Engineering & Remote Sensing.* 66(9):1113-1117.

Guth, P.L. 2004. The geometry of line-of-sight and weapons fan algorithms. *Studies in Military Geography and Geology.* Eds. Caldwell, D.R., Ehlen, J., and Harmon, R.S. Boston: Kluwer Academic Publishers. 271-285.

Hengl, T. 2006. Finding the right pixel size. *Computers & Geosciences.* 32:1283-1298.

Hodgson, M.E., Jensen, J.R., Schmit, L., Schill, S., and Davis, B. 2003. An evaluation of LIDAR and IFSAR-derived digital elevation models in leaf-on conditions with USGS Level 1 and Level 2 DEMs. *Remote Sensing of Environment.* 84:295-308.

Hodgson, M.E. and Bresnahan, P. 2004. Accuracy of airborne LIDAR-derived elevation: empirical assessment and error budget. *Photogrammetric Engineering & Remote Sensing.* 70(3):331-339.

Huising, E.J. and Gomes Pereira, L.M. 1998. Errors and accuracy estimates of laser data acquired by various laser scanning systems for topographic applications. *ISPRS Journal of Photogrammetry & Remote Sensing.* 53:245-261.

Lam, N.S. 1983. Spatial interpolation methods: a review. *The American Cartographer.* 10(2):129-149.

Lemmens, M. 2007. Airborne LiDAR sensors. *GIM International.* 21(2):24-27.

Li, Z. 1992. Variation of the accuracy of digital terrain models with sample interval. *Photogrammetric Record.* 14(79):113-128.

Li, Z. 1994. A comparative study of the accuracy of digital terrain models based on various data models. *ISPRS Journal of Photogrammetry and Remote Sensing.* 49(1):2-11.

Liu, X. 2008. Airborne LiDAR for DEM generation: some critical issues. *Progress in Physical Geography*. 32(1):31-49.

Liu, X., Zhang, Z., Peterson, J. and Chandra, S. 2007. LiDAR-derived high quality ground control information and DEM for image orthorectification. *GeoInformatica.* 11:37-53.

Llobera, M. 2007. Modeling visibility through vegetation. *International Journal of Geographical Information Science.* 21(7):799-810.

Lloyd, C.D. and Atkinson, P.M. 2002. Deriving DSMs from LiDAR data with kriging. *International Journal of Remote Sensing.* 23(12):2519-2524.

McCullagh, M.J. 1988. Terrain and surface modeling systems: theory and practice. *Photogrammetric Record*. 12:747-779.

Maune, D.F., Kopp, S.M., Crawford, C.A., and Zervas, C.E. 2007. Digital Elevation Models. *Digital Elevation Model Technologies and Applications: The DEM Users Manual.* Ed. Maune, D.F. Bethesda, MD: The American Society for Photogrammetry and Remote Sensing. 1-6.

Molander, C.W. 2007. Photogrammetry. *Digital Elevation Model Technologies and Applications: The DEM Users Manual.* Ed. Maune, D.F. Bethesda, MD: The American Society for Photogrammetry and Remote Sensing. 121-142.

Paine, D.P. and Kiser, J.D. 2003. *Aerial Photography and Image Interpretation.* Hoboken, NJ: John Wiley & Sons, Inc.

Riggs, P.D. and Dean, D.J. 2007. An investigation into the causes of errors and inconsistencies in predicted viewsheds. *Transactions in GIS.* 11(2):175-196.

Rose, S.M. 2001. *The effect of digital elevation model resolution on wave propagation predictions at 24Ghz.* Blacksburg, Va: University Libraries, Virginia Polytechnic Institute and State University. http://scholar.lib.vt.edu/theses/available/etd-05082001-130955.

Ruiz, M.O. 1995. *A Model of Error Propagation from Digital Elevation Models to Viewsheds.* Unpublished Ph.D. Dissertation. Department of Geography, University of Florida, Gainesville.

Ruiz, M.O. 1997. A causal analysis of error in viewsheds from USGS digital elevation models. *Transactions in GIS.* 2(1):85-94.

Smith, S.L., Holland, D.A., and Longley, P.A. 2003. The effect of changing grid size in the creation of laser scanner digital surface models. *Proceedings on the 7th International Conference on Geocomputation.*

Spencer, R. 2002. *Modeling and measuring GIS viewshed processing speed on different digital elevation model resolutions for line of sight radio wave propagation.* Blacksburg, Va: University Libraries, Virginia Polytechnic Institute and State University.

Su, J. and Bork, E. 2006. Influence of vegetation, slope, and LIDAR sampling angle on DEM accuracy. *Photogrammetric Engineering & Remote Sensing.* 72(11):1265-1274.

Virginia Information Technologies Agency. 2010. VBMP Orthophotography. http://www.vita.virginia.gov/isp/default.aspx?id=8412.

Xtreme Visual Basic Talk. 2006. Random Numbers. http://www.xtremevbtalk.com/showthread.php?=76270.

Yoeli, P. 1985. The making of intervisibility maps with computer and plotter. *Cartographica.* 22(3):88-103.