Agent-Based Simulation Optimization for Infrastructure Project Portfolio Selection

---

A Thesis

Presented to
the faculty of the School of Engineering and Applied Science
University of Virginia

---

in partial fulfillment
of the requirements for the degree

Master of Science

by

Thomas Ross Warner

December

2012

APPROVAL SHEET

The thesis
is submitted in partial fulfillment of the requirements
for the degree of
Master of Science

Tom Warner
AUTHOR

The thesis has been read and approved by the examining committee:

Dr. Donald Brown
Advisor

Dr. Matthew Gerber

Dr. Gerard Learmonth, Sr.

Accepted for the School of Engineering and Applied Science:

James H. Aylor
Dean, School of Engineering and Applied Science

December
2012

Table of Contents

Table of Figures

1. Abstract

This research concerns the geospatial optimization of portfolios of infrastructure projects within an urban context. Optimizing portfolios of infrastructure projects is not a trivial task since we must consider multiple objectives and the potential for hidden interactions that may exist between projects. To this end, this research utilizes an agent-based simulation called the Course of Action Planner (CAP) to capture the emergent outcomes of complex infrastructure systems with respect to a set of metrics. The CAP draws on geographic and demographic data to build an accurate representation of a specific urban environment. The environment, paired with simple rules for agent behavior and interaction form the foundation of the CAP simulation.

Building on Response Surface Methodology (RSM), this research offers a framework for designing experiments which will help determine local optimums for the locations of sets of infrastructure projects with respect to a given metric. Through a tractable example, this thesis tackles a portfolio optimization problem from start to finish. The example in this thesis seeks to find a location for a well in the city of Jalalabad, Afghanistan that minimizes the number of outpatient illnesses. Examples like this demonstrate the value of this approach in real-world applications. As a result, this research advocates a simulation optimization based approach to infrastructure project selection in situations where there are many human and social factors playing critical roles.

Existing frameworks for infrastructure prioritization and planning do not use simulation techniques which can capture social, economic and spatial phenomena that

can dictate a community’s response to various infrastructure projects. Optimizing simulations which contain both continuous and categorical variables requires a modification to be made to traditional RSM techniques. Generally, the goal of this research is the development of a framework that uses simulation optimization of an agent-based model to determine optimal portfolios and locations of infrastructure projects for a given city. While verification of results for such large systems remains a challenge, this thesis demonstrates an application and provides a solution.

2. Problem Statement

Infrastructure project selection is a problem that has faced governments, military strategists and global development organizations for centuries. The goals of typical infrastructure developments may range from improving traffic flow to increasing public safety, but all projects share an underlying goal of improving quality of life. Decision makers often lack quantitative means for comparing infrastructure projects across multiple and sometimes competing objectives. Under budgetary and time constraints leaders must be able to balance tradeoffs between projects in order to identify those which will create the best overall value-added for a system.

If ranking infrastructure projects against each other in terms of a single "value-added" metric were possible or even practical, then the problem of infrastructure project selection would be greatly simplified. However, the infrastructure selection problem is a classic example of a multi-objective optimization problem which requires a more complex decision process. Objectives may compete with each other; what represents an improvement in one metric may be a step backward for another. Decision-makers must have tools that can be used to weigh these tradeoffs to produce an optimal portfolio of infrastructure projects.

To make matters more challenging, there may also exist interactions between specific infrastructure projects which means that the combination of two or more projects may be worth more than the sum of the parts alone. Take the example of building a new road to connect a village to a new water source. The road alone may improve trade and transit and the water source alone may provide clean drinking water which is vital for public health. The synergistic effect of the two infrastructure projects is even greater than the sum of the two projects on their own. In essence, the road provides additional value to the water source. The lesson here is that there are complex interactions between infrastructure projects that cannot be ignored when evaluating the impacts of various portfolios of projects.

Although there are many examples of needs for infrastructure improvements, there is a specific case which will be the focus of this thesis. The United States Army Corps of Engineers (USACE) and the International Security and Assistance Force (ISAF) will be tasked with reconstructing Afghanistan in the next five to ten years. During this time, these organizations must restore ground security, establish self-governance and facilitate further development in one of the most war-torn nations in history. Afghanistan has been at the crossroads of ethnic and social conflict for the last 34 years. The country is in disarray and experts agree that the country will fall to ruin without intervention (Rashid, 2010). There is now a complete lack of infrastructure in a majority of the desolate and mountainous nation. After NATO forces have finished their battles in Afghanistan, there will be much to be done before they can leave Afghanistan. A coalition of forces, led by ISAF, must make genuine efforts to create a stable country, capable of ruling and controlling the affairs of its populace.

The United States has spent over 6 billion dollars on discretionary projects in Iraq and Afghanistan via Commander Emergency Response Program (CERP) funds since 2004 (Stein, 2011). While this is not the primary source of reconstruction funds, commanders believe that winning the hearts and minds of the people may be the key to counterinsurgency. According to General John Allen, Commander of ISAF, improving the quality of life of Afghanistan's citizens will create a safe, secure and non-threatening nation which is willing and eager to punish terrorists and warlords.

In order to stabilize the nation, many issues must be addressed by occupying forces. Counterinsurgency is a primary motivation for deployment of sustainable infrastructure in Afghanistan (Stein, 2011). Improving public health and safety is a top priority for ISAF. Hospitals, well networks and security forces are necessary to promote physical health and social well-being. The Afghan National Army (ANA) and the Afghan National Police (ANP) are the country's main security forces. They are being trained and supplied with resources by NATO forces. A sustainable, non-corrupt, security force is an essential facet of any development program.

Economically, the primarily agricultural country is also in shambles. With the opium trade fueling nearly one half of the nation's GDP (Gavrilis, 2010), significant changes in agricultural practices must be made in order to counteract future narcotics intervention. To assist with these changes, the United Nations Development Program (UNDP) is in the process of assigning an individualized "District Development Program" (DDP) to each of Afghanistan's 402 districts. These DDPs are designed to "enhance district-level governance to deliver services to the poor and vulnerable, as well as improve sustainable and diversified livelihoods through productive infrastructure." Commanders need a tool that can quickly and quantitatively compare portfolios of infrastructure projects to aid in the design and deployment of DDPs.

3. Literature Review

The literature review compiled herein aims to serve as a foundation for the methods developed in Section 4. It is imperative to begin this section by describing existing methodologies for infrastructure prioritization in developing nations. Specifically, this review will focus on health infrastructure planning, water infrastructure standards and the general need for security when implementing infrastructure projects effectively. Next, this review describes the notion of infrastructure projects as counterinsurgency.

The literature review then goes on to define the simulation optimization technique which can be applied in the context of infrastructure project portfolio selection. The next section documents the rationale behind agent-based simulation as one of the preeminent tools for modeling of humans and social behavior. The topic of spatial statistics is mentioned due to the geospatial nature of this particular genre of portfolio selection. Ultimately, this review provides the evidence and the reasoning behind the methods and conclusions derived in the following sections.

3.1. Existing Frameworks for Infrastructure Prioritization in Developing Nations

Lambert et al. (2012) uses a risk-informed multi-criteria analysis methodology for prioritizing infrastructure projects in Afghanistan. This approach utilizes a three layered method for infrastructure investment analysis. Layer I, known as Inter-Industry Performance, is tied to the provincial economic opportunities provided by an infrastructure project. Layer II, called Inter-Project Performance uses multiple stakeholders' perspectives to build a multi-criteria decision aiding (MCDA) model which scores the investment’s performance in comparison to alternative infrastructure projects. The Layer II analysis attempts to account for emergent conditions based on the impacts of an infrastructure project with respect to multiple stakeholders. Some of these emergent conditions include “increased regional trade” and “raw materials decrease.” These conditions can occur as the byproduct of certain combinations of infrastructure projects and are built into the MCDA model. The third and final Layer, called Project Performance and Requirements, is an analysis of the resources and scheduling needs for an infrastructure investment. The risk-informed multi criteria analysis proceeds by assigning weights to each infrastructure investment. By quantifying various aspects of the infrastructure projects, it is possible to determine optimal projects.

An essential aspect of this holistic approach is to have a stakeholder workshop. If it is possible to collect a group of stakeholders to include citizens, local leaders, and developers, then this approach to infrastructure prioritization can be very favorable. In many cases, however, it is simply not feasible to have a stakeholder workshop and it is therefore impractical to use this framework.

3.1.1. Health Infrastructure in Developing Nations

Unger & Criel (1995) describe some of the key principles of health infrastructure planning in less developed countries. It is useful to consider several of these principles when evaluating the impacts of a proposed health infrastructure project such as a hospital. The key points are that a health system should be entirely integrated at all levels and that intermediate health facilities are not an efficient use of resources. For example, say that an intermediate facility covers a population of 80,000 people. This facility actually winds up employing more staff (nurses, midwives, doctors, etc.) than eight smaller (serving 10,000 people each), decentralized facilities would. The intermediate facility would not be of the same scale as a hospital and therefore would not be able to provide emergency services like ambulances nor would it have the same level of technology available as a hospital. Hospitals are an essential piece of infrastructure for an urban region, but it is still important to offer primary care at local health centers which can be accessed on foot from anywhere in the city. Unger & Criel would argue that in order to complement an existing hospital in Jalalabad, health centers must be built throughout the city to provide primary care and to help foster public health awareness.

3.1.2. Water Supply Infrastructure

Domestic water supply is defined by the World Health Organization as all water which is used for consumption, bathing and food preparation. Access to clean water is a critical component of the United Nations Development Programme's (UNDP) Human Poverty Index for developing countries. The Sphere Project sets the minimum average domestic water demand at 15 liters per person per day. Additionally, The Sphere Handbook instructs that the maximum distance from any household to the nearest water point is 500 meters and that the queuing time at the water source is no more than 30 minutes. If the queuing time is too long or the water source is too far away, a person may resort to procuring water on their own. This can be a dangerous practice if care is not taken to ensure the water quality. This is especially true in an agricultural region since there will be certain chemicals present in water sources which make contact with the ground.

The handbook also states that even if there is a sufficient quantity of water available to meet minimum needs, there needs to be some additional measures put in place to ensure that all groups are afforded equitable access to the water source. In a nation like Afghanistan, it is important to consider the possibility that certain ethnic minorities or women may not receive equal access to water sources.

There are a plethora of risks associated with limited access to clean drinking water. Severe dehydration can be fatal. Dehydration will result from failure to consume enough water. It can also occur as a short-term effect of the loss of body fluids due to diarrhea. Chronic dehydration can lead to urinary stone formation. If access to safe drinking water is limited for any reason, a person may be compelled to find water from existing surface water sources. Fecal matter and other harmful bacteria can exist in surface water sources which can cause E. coli, cholera and a litany of other illnesses. Therefore, it is critical to the health of a population that a locality adheres to the guidelines laid out by The Sphere Project Handbook.

3.1.3. Security for Infrastructure Projects

An article by Moser & McIlwaine (2006) discusses a framework for urban violence reduction in Latin America which centers on urban infrastructure developments. The article establishes that crime has become ubiquitous in urban Latin America. Moser challenges the common stereotype that poverty is the primary cause of violence. Instead, the argument is that the inequality and exclusion associated with unequal distribution of economic, political and social resources in urban contexts is the real driver behind violence. Unequal access to or total restriction from physical infrastructure are often factors that contribute to crime. It is imperative that all citizens maintain equal access to public infrastructure projects. Although this article focuses in Latin America, the concept of cultural groups organizing against one and other hits home in Afghanistan. Instead of gangs, ethnic groups control parts of cities and currently have control over who receives access to critical pieces of infrastructure.

In August 2012, an Afghan newspaper, the Khaama Press, reported that the Afghan President, Hamid Karzai, had ordered security organizations to double their measures for all infrastructure projects in the country due to a rise in insurgency and an “unstable situation in various regions of Afghanistan”. Highway projects, electricity projects, and fiber optics communication projects have been facing the most security threats. 53 construction workers have been killed by insurgents since 2005 and many projects have been delayed or suspended due to security threats. Infrastructure projects often become the target of insurgent activity because they represent development and therefore a threat to the insurgency.

3.1.4. Infrastructure as Counterinsurgency in Afghanistan

U.S. Army Field Manual 3-24, titled Counterinsurgency, outlines the strategic approach being taken to combat insurgency. The manual covers all of the basics of counterinsurgency (COIN): civilian participation in military operations, intelligence, design and execution of COIN operations, building host-nation security forces and sustainment. The field manual advocates a “Shape-Clear-Hold-Build” approach towards incremental infrastructure improvement. Building host-nation security forces is seen as a critical step in COIN operations since the goal is to provide for long-term regional stability. Providing basic health based infrastructure projects can aid COIN operations as quality of life improvements are a threat to insurgent forces.

Building infrastructure projects to aid in counterinsurgency was one of the operations credited with turning the tide in the Iraq War. Unfortunately for NATO and global development organizations, the same cannot be said for the War in Afghanistan (Reuters, 2012). According to the Special Inspector General for Afghanistan Reconstruction (SIGAR) almost $400 million spent in 2011 on power grid, highway and other infrastructure projects have not achieved their desired counterinsurgency effects. The SIGAR went on to add, “In some instances, these projects may result in adverse COIN effects because they create an expectations gap among the affected population or lack citizen support.” Although, they may not be seen for years to come, the US Department of Defense remains optimistic that there will be positive counterinsurgency effects from the current spending in Afghanistan (Reuters, 2012).

3.2. Response Surface Methodology for Simulation Optimization

Traditional response surface methodology (RSM) was first introduced in 1951 by Box & Wilson to serve as an efficient means for optimizing a specific response of a given process. RSM has been used to optimize a variety of processes ranging from physical and chemical processes (Vincente, et al., 1998) to computer simulation models (Montgomery & Bettencourt Jr., 1977). When dealing exclusively with continuous independent variables, it is possible to use RSM to attain a local optimum. RSM is based on sequential experimentation. Each experiment can be thought of as a design point on the path to an optimum (Montgomery, 2008).

The experimental region of interest is defined as the range of possible values for each independent variable in the process. The first step in RSM is to select a starting point in this region of interest. This point will serve as the center for a factorial design of experiments. It is encouraged, although not necessary, to run multiple replications at each design point. By performing *m* replications at each design point in a factorial design with *n* independent variables, an experimenter will need a total of $m*2^n$ experiments.

A first-order model is fit to the data and the direction of steepest ascent (or descent in a minimization problem) is created. Beginning at the initial starting point, proceed outwards into the region of interest along the direction of steepest ascent. Next, perform experiments at points on this path in step sizes determined by the experimenter (based on process knowledge and other practical considerations). Continue choosing design points along this direction until no further increase in the response is found. At this point, the entire process can begin again, or, if the experimenter is satisfied with the region of the new optimum point, a second order model can be fit based on a Central Composite Design (CCD).

CCD requires design points that exist on a sphere centered on the optimum point found from the steepest ascent. CCD designs consist of another 2^n factorial runs, 2n axial runs and center runs. This data will have the property of rotatability which means that all runs are the same distance from the center point of the design. The variance of the predicted response will be constant on spheres which is useful when trying to locate a local optimum.

The concept of using meta-models to approximate simulation behavior over a region is taken from Kleijnen (2009). The author talks about using low-order polynomials which are linear regression meta-models to predict simulation behavior over small regions where CCD sampling has taken place.

In the case of infrastructure project portfolio optimization, not all of the independent variables are continuous. The class of infrastructure project is a critical piece of the puzzle and there must be some way to account for this within the RSM framework. Lenth (2009) introduces an approach to RSM with categorical variables:

> “In practice, categorical variables must be handled separately by comparing our best operating conditions with respect to the quantitative variables across different combinations of the categorical ones.”

This means that infrastructure project portfolios can be optimized only when the numbers of infrastructure projects in each class are fixed. In order to compare portfolios with different project compositions, the concept of Pareto-optimality must be considered.

3.3. Agent-Based Simulation for Social Modeling of Human Populations

In order to address the issue of infrastructure project scoring, I propose the use of agent-based modeling due to its documented uses with complex systems. Agent-based modeling has been used to uncover hidden interactions and capture the essence of emergent outcomes in complex systems. Literature has shown the importance of agent-based modeling in four fields of simulation: flow, organizational, market and diffusion (Bonabeau, 2002). Agent-based approaches are a way of interpreting a problem rather than a specific technology. Agent-based modeling has been catching on as one of the premier methods of social modeling. The three primary benefits of agent-based modeling are that it can capture emergent phenomena, that it provides a natural description of a system and that it is highly flexible. By formulating a simulation environment with agents, this research aims to compare and score the effectiveness of infrastructure project portfolios. This model will help to uncover any emergent phenomena which may be a product of interactions between individual infrastructure projects.

As Moser & McIlwaine (2006) has shown, inequities in access to physical infrastructure systems can be a source of violence. It is therefore critical to model agent-level ethnic differences. This level of detail will model the struggle that minorities can face when they try to extract a resource (water) or utilize a service (hospitals). Priority is given to ethnic majorities and to wealthy individuals who can literally buy their way to the front of lines and receive priority access resources. Agent-based simulation is an effective way to portray the impacts of these individual differences between citizens on the way infrastructure is accessed in the model.

3.4. Geospatial Statistics

Due to the highly geospatial nature of portfolio selection, it is imperative that appropriate statistics for spatial data are applied. The location of hazardous events and crimes in general may be dictated by the locations of various infrastructure projects. Other work describes spatial point patterns as collections of random events (Cressie, 1993). These realizations can be used to infer parameters of the point process. Generally speaking, if we can model the occurrence of these security breaches, we can develop a threat map which can be used to improve security strategies, an integral part of infrastructure development. The same thought applies to the deployment of health services, such as hospitals or wells for clean drinking water, to match the point pattern of sicknesses. The RSM approach for simulation optimization will be used to find the optimal site locations for specific pieces of infrastructure.

4. Methods

To address the problem of portfolio selection, this research utilizes a multi-tiered approach. This approach involves optimizing an infrastructure project portfolio by using the RSM approach with an agent-based. The agent-based model will be herein referred to as the Course of Action Planner (CAP). CAP involves the use of Geographical Information Systems (GIS) to characterize the environment and response surface methodology to optimize the location of specific infrastructure projects within a portfolio. Unlike most existing decision support tools, CAP will provide an analytical, quantitative means for optimizing portfolio performance with respect to a specific metric in a specific environment given a specific budget.

4.1. Agent-Based Simulation

Recent work provides us with a three-tiered definition of the components of an agent-based model (Macal & North, 2010):

1. A set of agents, their attributes and behaviors.
2. A set of agent relationships and methods of interaction: An underlying topology of connectedness defines how and with whom agents interact.
3. The agents' environment: Agents interact with their environment in addition to other agents.

Extending this definition to our specific case, agents will represent individual citizens of the city. Their attributes include their level of wealth, ethnicity, current location, house location, and health status. Agents interact with each other and their environment according to a set of rules related to their current health status.

Agents' health is modeled as an integer between 1 and 4 inclusive. A health of 1 corresponds to a healthy agent with no need for a hospital. Agents with a health of 2 required an outpatient visit to the hospital. Health scores of 3 are given to agents with a need for urgent care. The final health score possible, 4, was used to denote a deceased agent.

4.1.1. Application

Picking the potential locations for infrastructure projects is only half of the battle. The second function of CAP is to determine how each different combination of projects, or portfolio, will affect the population. The CAP simulation was built using the parameters of a sample city in Afghanistan. Jalalabad, in northeastern Afghanistan is used as a proof-of-concept because of its relatively large population of 170,000 people. Jalalabad was also selected due to its relative strategic importance. Jalalabad is the home of US military base and along the major Asian highway connecting Kabul to the Peshawar province in northern Pakistan. After the possible infrastructure project sites have been selected, an agent-based simulation will provide a quantitative comparison of the effects of the portfolio of infrastructure developments.

4.1.2. Data

Geographic data provided by the National Geospatial-Intelligence Agency (NGA) was used to build the environment. This data included the location of government buildings, roads, residential areas, farmland, water, forested land and "open" area. Demographic data had to be compiled through various sources to get an accurate representation of the population's composition. Life expectancy, sickness rates and ethnic distribution data were collected from the World Bank, CIA World Factbook, DDPs and the 2004 Afghanistan National Hospital Survey Report. Table 1 summarizes the key data that were used in the development of CAP.

Table 1 Selected Data and Sources

Afghanistan Data	Value	Source
Life Expectancy	48.3	World Bank, 2010
Poverty Rate	36%	World Bank, 2008
Rural Access to Improved Water	42%	World Bank, 2010
Jalalabad Data	**Value**	**Source**
Annual Instances of Outpatient Sickness	4,792	ANHSR, 2004
Population	170,000	Encyclopedia Britannica, 2006
Pashtun %	85%	UNHCR Sub-Office Jalalabad District Profile, 2002

4.1.3. Implementation

CAP's agent-based simulation was coded in Repast Simphony 1.2 based on the RepastCity project (Malleson, 2011). Figure 1 shows an example of a simulation in Repast Simphony 1.2 with agents (tiny orange squares), roads, houses (grey pentagons), hospitals (green triangles) and wells (blue circles). 200 agents are used to model Jalalabad's 170,000 people. Agents' behavior follows rules based on the status of their health which is affected by their access to resources and encounters with danger. The simulation keeps track of the number of health related incidents that occur over the course of a two month period. These metrics are used to compare the differences between the various portfolios of projects.

Figure 1 Screenshot of a simulation in Repast Simphony 1.2 showing agents, roads, wells, and hospitals.

4.1.4. Participatory Modeling

Although it was not used in this research, it is highly recommended that stakeholder feedback be involved in the agent-based modeling process. This validation step can help define the relationships that are used in the logic that underlies the model. Critical members of this discussion would include infrastructure developers, political leaders, citizens who would be affected by the projects, and members of NGOs that serve in developing nations. Everyone would bring a unique perspective to the model building effort and this would help to uncover some relationships that could otherwise be overlooked. Becu, et al. (2008) discusses the limitations of participatory computer simulations to support collective decision-making. Among the complaints, Becu notes that it was difficult to get a group of farmers to understand that the simulation was to model a reproduction of reality and not to model reality itself. To this end, although companion modeling can be very useful, it can be difficult to get the right kind of help from all stakeholders involved.

Barreteau et al. (2001) and Barreteau, Bousquet & Attonaty (2003) offer examples of companion modeling being used successfully to improve the quality of agent-based simulations. In 2001, Barreteau et al. showed how collaboration with farmers in the Senegal River valley was used to help open the black-box that had been used to describe multi-agent systems. Role playing games were used to determine how humans would react to various scenarios in their research. In 2003, Barreteau, Bousquet & Attonaty used role playing games to show how artificial societies could be created based on stakeholder feedback. If the concerns of Becu et al. are considered and handled appropriately, the literature has shown several instances of participatory modeling which have led to improved agent-based models.

4.2. Response Surface Methodology for Geospatial Optimization

CAP can be optimized to prescribe the best geospatial orientation of infrastructure projects with respect to an objective function or a counterinsurgency performance metric. To further address the problem of infrastructure site selection, RSM will be used. In this case, CAP can be used to collect data to determine optimal well and hospital locations with respect to a health metric. For instance, one orientation of hospitals and wells may yield a significant decrease in deaths in the simulation, while another may not show any improvement over the status quo. These types of results will give infrastructure planners vital strategic planning insight.

4.2.1. Categorical Variables and RSM Optimization

The presence of categorical variables, namely the different classes of infrastructure projects, makes using a general RSM approach, like the one outlined above, infeasible. It becomes necessary to determine how many of each class of infrastructure projects are needed. For example, consider the following situation:

1. Only wells and hospitals are being considered.
2. Between one and three wells must be built.
3. Between zero and two hospitals must be built.
4. The goal is to minimize the expected number of deaths in a one year period after construction.

It is easy to see that there are nine possible combinations of infrastructure projects that can be considered. Each of the nine combinations has a total cost estimate based on the quantities and prices of each type of project in the combination. For each of these nine combinations, the aforementioned RSM approach can be used to find the optimal project orientations. This is done by treating the site location of each infrastructure project, *i*, as an ordered pair (x_i, y_i) on a continuous scale over the geographic scope of the city.

4.3. Pareto-Optimal Portfolios

The optimal project orientations for the nine combinations will give a decision-maker a quick reference to the tradeoffs between cost of a combination and expected number of deaths given the optimal site locations. By examining one's budget and the Pareto-optimal portfolios, any decision-maker can make a sound, analytical decision.

Table 2 shows the data that is plotted in Figure 2. Pareto-optimal portfolios are those which are not dominated by any other portfolio in terms of cost and deaths. Charts like Figure 2 can be used to visualize Pareto-optimal portfolios. This is an end product that a decision-maker would be interested in to determine which portfolio to select.

Table 2 Data from optimal configurations of all combinations of hospitals and wells. Pareto-optimal portfolios are listed in italics.

Wells	Hospitals	Cost ($M)	Deaths
1	*0*	*200*	*16*
1	1	700	14
1	2	1200	10
2	*0*	*400*	*11*
2	1	900	10
2	2	1400	7
3	*0*	*600*	*9*
3	*1*	*1100*	*7*
3	*2*	*1600*	*6*

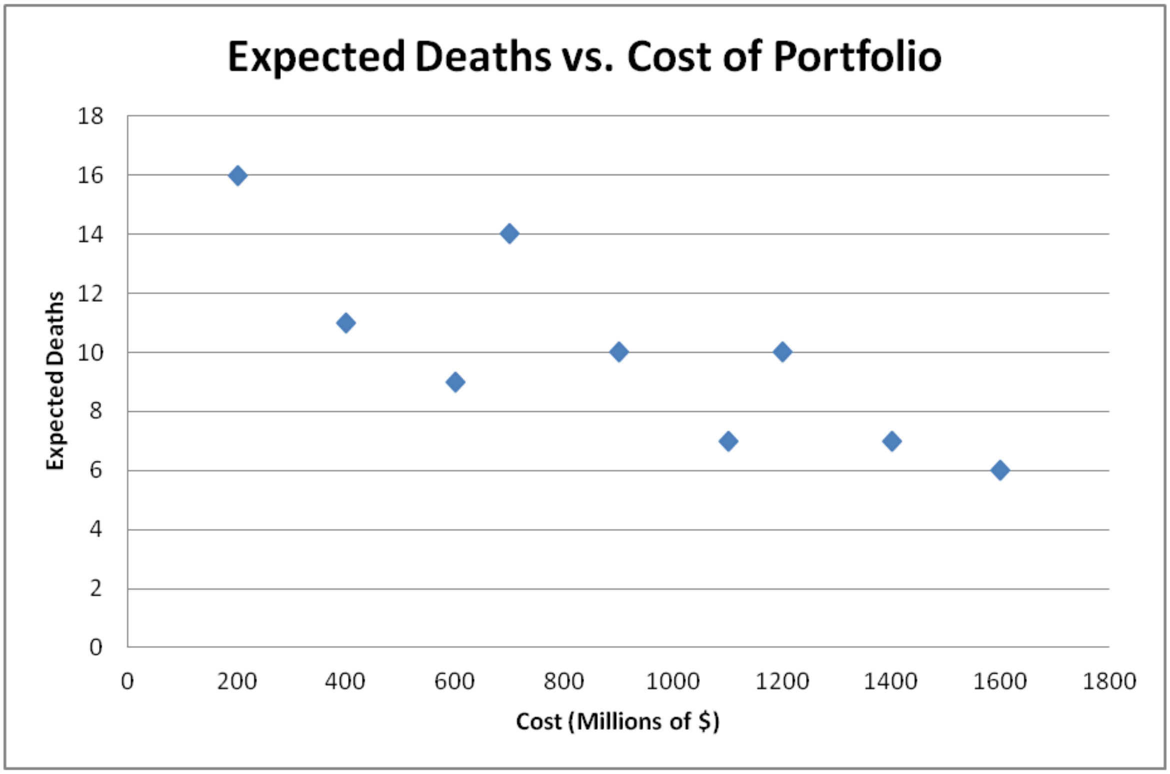

Figure 2 Cost of portfolio vs. deaths.

5. Results

This section of the paper will chronicle the process of obtaining an optimal site location for a well project in the Jalalabad environment with respect to minimizing the number of outpatient illnesses. This is a practical example because the primary cause of outpatient illness in Afghanistan is due to a lack of clean drinking water and proper sanitation practices. An optimal well location will minimize the total number of outpatient visits to hospitals in the simulation. Literature has shown that proximity to a water source has a direct impact on improvements in citizens' health. Since RSM will only yield a local optimum, it is important to choose a logical starting point and a reasonable scale for the initial factorial design. If the initial starting point is not practical, it is possible that the local optimum found will not be a useful solution.

5.1. Factorial Design

The locations of each of the factorial design points should be determined by choosing an appropriate center point and axial distance. Points for each infrastructure project should be kept as two separate coordinate values. It is important to maintain consistent distance measures so that the direction of steepest ascent (or descent) can be ascertained from the data collected at the factorial design points. If the simulation is modeling a stochastic process and resources permit, it is best to take multiple replications or runs at each of the design points. This will help to refine the direction of steepest ascent. For this example, three replicates were measured at each design point for a total of twelve runs. In order to assign these variables, one must break up the region into a grid.

From this point forward, the region of interest is modeled as a 4x4 grid. The initial factorial design is shown in Figure 3 and is centered around the point (1,1). All of the design points can be seen in Appendix A.

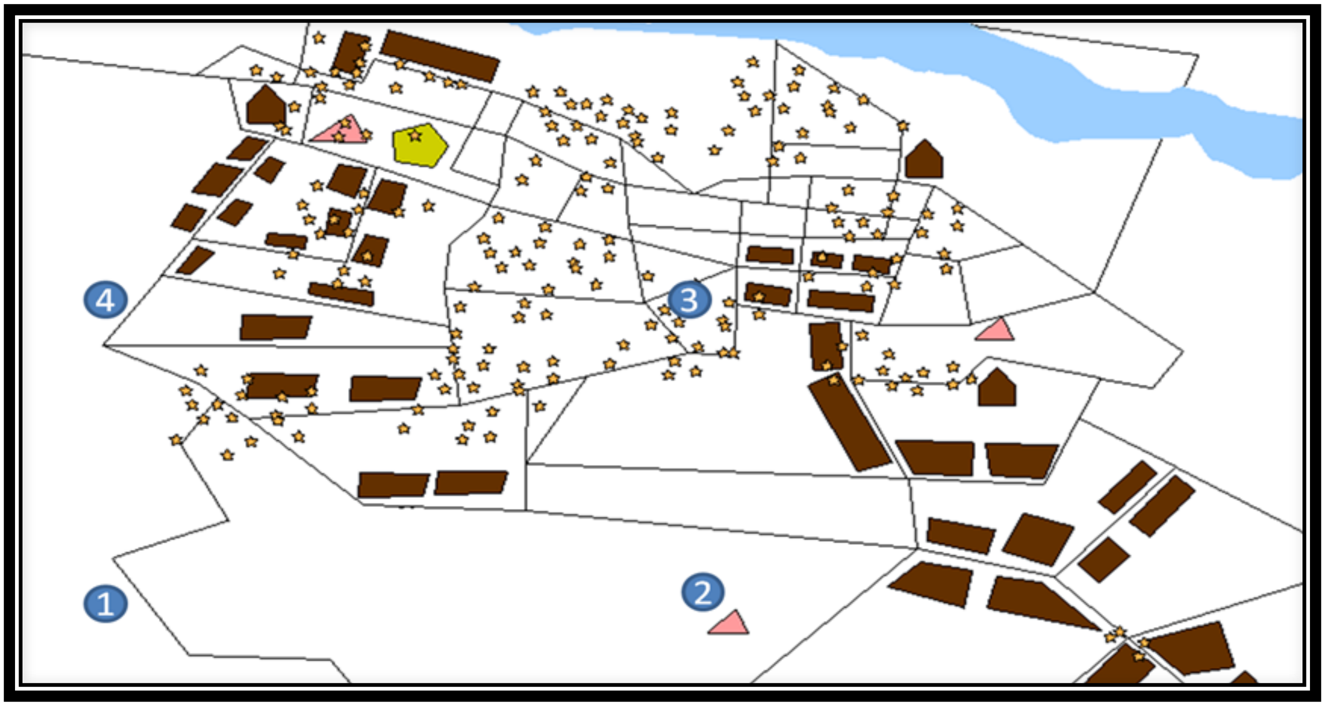

Figure 3 Screenshot of the CAP simulation showing the four well locations tested in the factorial design.

5.2. Steepest Descent Direction

The steepest descent direction is found by fitting a first order linear model to the data collected in the factorial design. This direction is shown visually by the plane of best fit plotted with the 3D scatterplot in Figure 4.

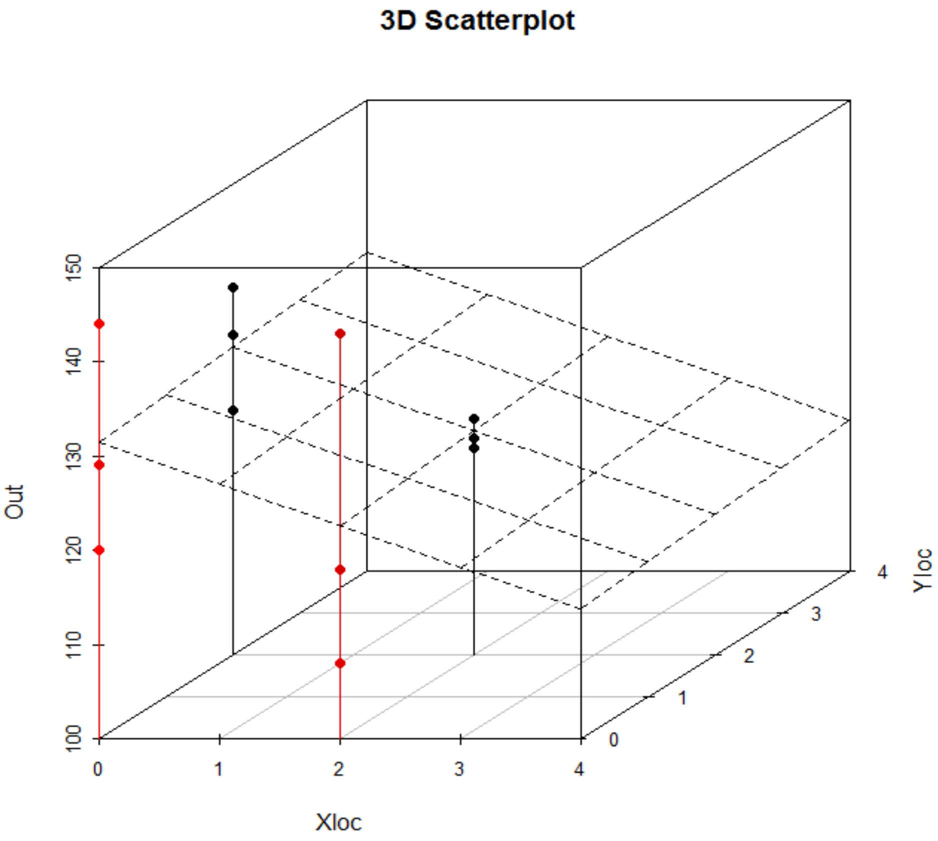

Figure 4 3D Scatterplot depicting the direction of steepest descent.

After collecting three replicates at each point, a linear regression was fit to the data. The linear regression model is shown in Appendix B. The direction of steepest descent was found to be <0.99, -0.13>. The next series of design points for the simulation are found by proceeding out in this direction in discrete steps. Again, replicates should be taken at each step along the way in order to account for standard variance in the simulation. The map in Figure 5 shows this new design point series as the points (5, 6, 7, 8, and 9). They are shown in a table in Appendix A as well.

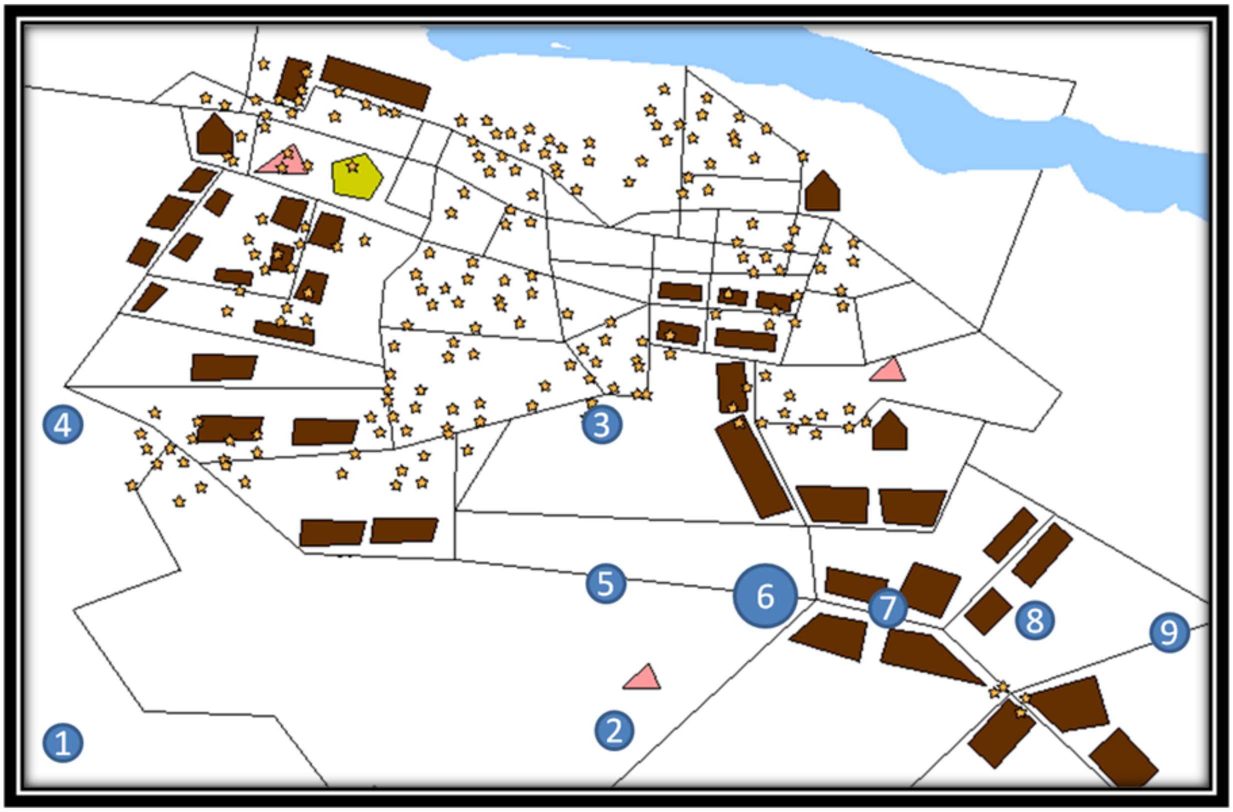

Figure 5 Screenshot of the CAP simulation showing the five well locations tested in the steepest descent direction (5-9) as well as the original factorial design points (1-4).

Continuing out on this path, a minimum is found at the second step (the first step was on the face of the factorial design region). This is shown in the steepest descent plot below. The mean number of outpatients given this well site was only 112.33 which was a significant improvement over the previous designs that had been tested. This improvement leads us to believe that we are near an optimal well site location at (2.5, 0.8). Given this site's location at an intersection between two roads, it is logical that this site would be well-suited for a high-traffic well.

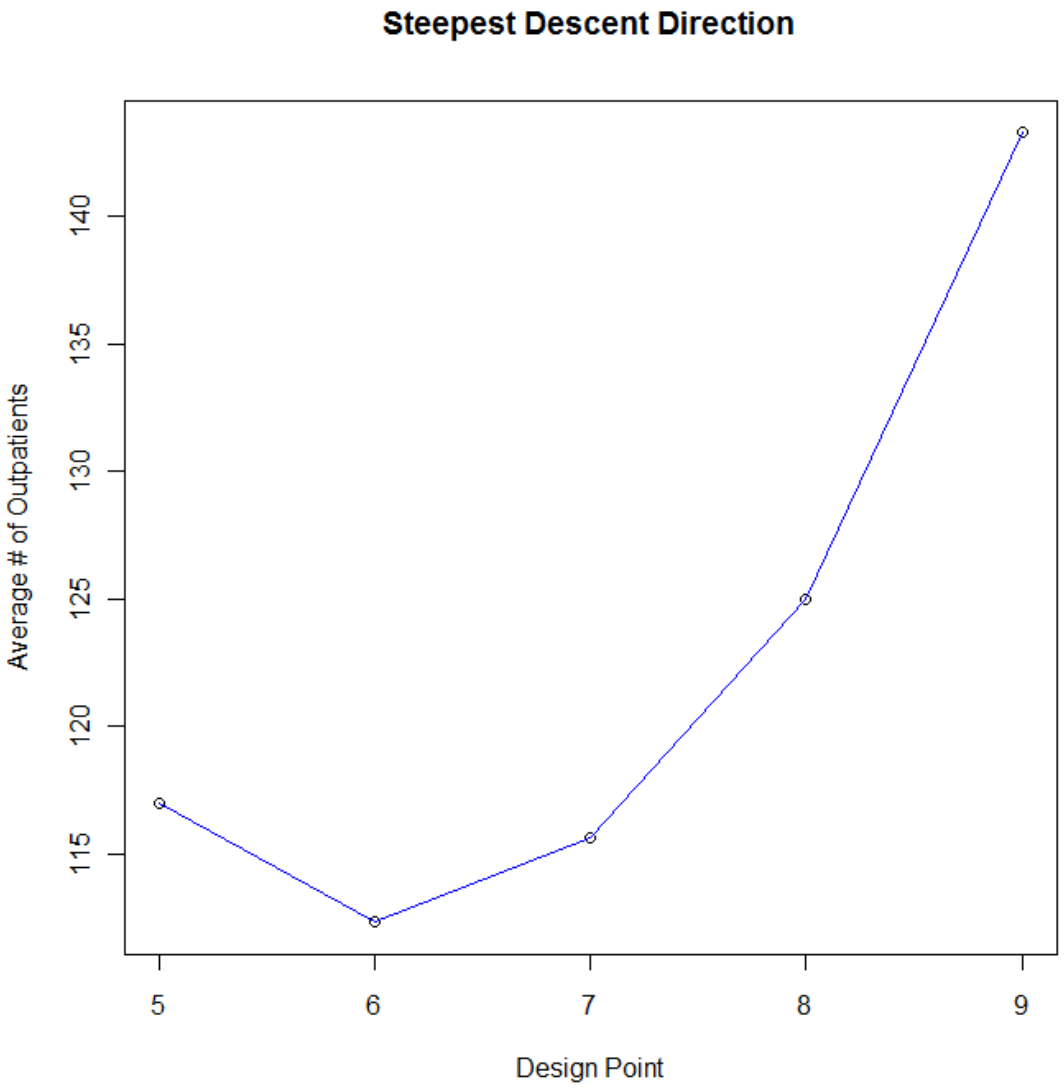

Figure 6 Averages of the number of outpatient visits from the five “steepest descent” design points.

As Figure 6 shows, the minimum average number of outpatients occurs at design point number 6. This point is then selected as the center point for a Central Composite Design (CCD).

5.3. CCD

For this example, only one project is being located so only two factors are being examined. In this case, eight new design points must be utilized in addition to the center point making nine total points. In general, the number of points needed to run a CCD is equal to $(2^k + 2k + 1)$ where k is the number of factors being simulated. Since every infrastructure project has two factors (latitude and longitude), the number of CCD points will be $(2^{2n} + 4n + 1)$ where n is the number of infrastructure projects. Unfortunately, this number will grow very quickly with the number of infrastructure projects as shown in Table 3.

Table 3 The number of CCD points grows exponentially with the number of projects in a portfolio

Projects	Factors	Factorial Points	Axial Points	Total CCD Points
1	2	4	4	9
2	4	16	8	25
3	6	64	12	77
4	8	256	16	273

The data used for the CCD are in Appendix A and the higher-order model that was fit is shown in Appendix B. Since the p-value of the CCD model is 0.31, which is greater than 0.05, we must fail to reject the hypothesis that at least one of the coefficients in the model is significantly different than zero. This means that no significant higher order model can be used to approximate the behavior of the response surface around the center point (2.5, 0.8). In this case, it is wise to choose the point which had the lowest average number of outpatient visits from the CCD points tested. This point happens to be number 14 at (2.61, 08) where an average of 94.67 outpatient hospital visits were recorded. This number represents a marked improvement from the center point of the CCD and any of the points tested in the initial factorial design.

5.3.1. Rotatability and a Higher-Order Model

CCD’s have the nice property of being rotatable which means that all design points are the same “distance” from the center point and therefore the variance will be constant at each point. It is trivial, then, to build a higher-order model around the center point. We can assume that the higher-order model will be a reasonable approximation to the simulation’s behavior around the center point. By setting the model’s derivative to zero, it is possible to find the optimal site location for the well. This process is shown below. All design points can be found in Appendix A. All summary statistics and linear models can be found in Appendix B.

5.4. Verification

Since there is no way to verify that a site location is truly optimal, the best verification can be done by visual inspection. Once the infrastructure is deployed, it is instructive that the metrics of interest, in this case, outpatient visits to hospitals in the city, are kept to document the improvement. If it is possible to include local officials, citizens and other stakeholders in the decision process, then each group will also be able to provide some feedback.

6. Conclusions

6.1. Benefits of Approach

There are some key benefits that this approach offers. Simulation offers researchers of infrastructure development an opportunity to build an environment and test hypotheses without having to physically implement systems. The entire process of building an agent-based model is modular. The more that is known about a specific relationship or phenomenon, the better it can be modeled. In the event that very little is known about a process, it is possible to black box it, meaning that only the inputs and outputs need to be known and not the underlying relationships between them. The more accurate the simulation is, the more useful the optimization.

The second main benefit of this approach is that the experiments are repeatable. Repeatable experiments mean that we can get a variance for each data point which allows us to build confidence around our recommendations. As with any kind of stochastic modeling, it is very useful to know the variance of a model. If the regressions fit to describe the simulation's first-order behavior are not significant, then it is impractical to go forward with the RSM approach.

The third major benefit of this method is that the entire process can be generalized to a variety of situations. It is simple to change the map, thereby changing the road network, the houses, the hospitals, the wells and other features. The metrics don't necessarily have to be health related. It is possible to model economic processes, traffic operations and other human & social behaviors (Bonabeau, 2002). The geospatial RSM approach outlined in this thesis could be applied to a wide range of applications involving the optimization of multiple points on a grid. Extensions of this work could include disaster relief planning to answer the question: where would critical resources be stored and distributed most effectively?

6.2. Drawbacks of Approach

The drawbacks of agent-based simulation optimization for infrastructure project portfolio selection were primarily in terms of validation, verification and complexity. This particular case study in Jalalabad presented some unique challenges, but also served as a good opportunity to go where no other simulation has gone before. For this reason, it is essential to highlight a few of the drawbacks that were experienced with this approach.

6.2.1. Validation

In order to use this specific approach to modeling, one must acknowledge that the results found by optimizing the CAP simulation provide only a theoretical local optimum. That is, the results are only as good as the assumptions which underlie the model. The more that is known about the interactions in the simulation, the more accurate the results will be. For this reason, it is critical that some form of participatory modeling take place prior to implementation. The simulation built for this thesis relied on input from an Army Major and a couple contact points at the Army Research Lab in Aberdeen, MD. Ideally, there would be several more key stakeholders involved in the simulation building process.

6.2.2. Verification

Verification of this infrastructure planning technique is a challenging issue. The verification issue stems from the fact that 200 agents were used to model a population of 170,000. The number of sicknesses, critical injuries and deaths must be computed on a much higher level of granularity than this simulation affords (1:850). The actual number of each health incident is not as important as the relative change between different infrastructure project portfolios. One solution to this problem is to normalize all of the results so that the number of deaths in the “status quo” simulation is the same as the actual number of deaths expected in a given time frame. If the relative percentage increase or decrease in a metric is accurate, then it is reasonable to use this scaling approach. If this approach does not appear to hold for a given scenario, then alternative verification methods must be explored. It is critical to note that the methods, results and conclusions highlighted in this thesis are just one example of how agent-based simulation optimization can be used for infrastructure project portfolio selection. The technique of modeling and optimizing portfolios of projects can and should be generalized for other applications. It is conceivable that verification would make more sense in the context of other case studies, although modeling complex phenomena almost always requires some form of empirical verification, which is simply not attainable for this example.

6.2.3. Complexity

The RSM approach requires several iterations of simulations to be run. As the number of projects in a portfolio increases, the number of simulations increases exponentially. By the 4th project, a total of 273 design points must be tested just for the

CCD alone. If the design points could be automated, then this would not be a highly critical issue for portfolios with five projects or less. Unfortunately, in this framework, each design point requires a unique map to be applied. The process of adjusting a map for a given design point takes a few minutes of an operator's time and will make optimization very tedious for portfolios with more than two infrastructure projects due to the number of CCD design points as shown in Table 3. Automation of the RSM design points is an area of work that should be considered if this technique is to be used extensively in the future.

7. Future Work

7.1. Automation of RSM

There is still a great deal of future work that can be done to improve this experimental design framework. In order to tackle the limitation of complexity due to the RSM technique, it may be possible to automate the creation of the maps. If map creation can be automated, then it would be possible to dynamically run the simulations prescribed by the RSM approach. Currently, infrastructure projects must be edited manually in ArcGIS or similar map editing software. Due to the nature of factorial designs and CCDs, the number of maps required grows very quickly with the number of infrastructure projects being used. A full automation of the RSM search for an optimum could speed the total simulation optimization time up significantly.

7.2. Data Validation and Verification

Enabling stakeholder feedback in the simulation development process would help refine some of the rules which define the interactions between agents and infrastructure. In terms of data validation, lots of work can be done to improve the quality of the data entering the simulation. As far as the outputs of the simulation go, it is impossible to get empirical data that supports or refutes the findings of the CAP simulations. This is because it is not practical to actually install a given portfolio of infrastructure projects. If there were a situation where an infrastructure project was being built, it would be extremely beneficial to record the changes in some related metrics. This would give a sense of the accuracy of the simulation and would allow the model to be refined or rescaled. This would be a critical step in verifying the agent-based modeling approach for infrastructure project prioritization.

References

Rashid, M. Taliban: Militant Islam, oil and fundamentalism in Central Asia, Second Edition, New Haven, CT: Yale University Press, 2010.

Stein, M. Lawmakers question CERP funds in Afghanistan. Army Times, August 4, 2011.

Gavrilis, G. The good and bad news about Afghan opium. Council on Foreign Relations, February 10, 2010.

Box G. & Wilson, K. On experimental attainment of optimum conditions Journal of the Royal Statistical Society, no. 13, pp. 1-45, 1951.

Vicente, G., Coteron, A., Martinez, M. & Aracil, J. Application of the factorial design of experiments and response surface methodology to optimize biodiesel production, Industrial Crops and Products, Volume 8, Issue 1, March 1998, Pages 29-35, ISSN 0926-6690, 10.1016/S0926-6690(97)10003-6.

Montgomery, D.C. & Bettencourt Jr., V.M. Multiple response surface methods in computer simulation. Journal of Simulation. October 1977 vol. 29 no. 4 113-121. doi: 10.1177/003754977702900406

Montgomery, D.C., 2008. Design and Analysis of Experiments, 7th edition. New York, NY. John Wiley & Sons, Inc., 1993.

Kleijnen, J.P.C., 2007. Design and Analysis of Simulation Experiments, (1st ed.). Springer Publishing Company, Incorporated.

Lenth, R. V., 2009. Response-surface methods in r, using rsm. Journal of Statistical Software, 32(7):1-17. Retrieved at http://www.jstatsoft.org/v32/i07/paper

Lambert, J., Karvetski, C., Spencer, D., Sotirin, B., Liberi, D., Zaghloul, H., Koogler, J., Hunter, S., Goran, W., Ditmer, R., and Linkov, I. Prioritizing infrastructure investments in Afghanistan with multiagency stakeholders and deep uncertainty of emergent conditions. Journal of Infrastructure Systems 2012 18:2, 155-166.

Unger, J.-P. and Criel, B. (1995), Principles of health infrastructure planning in less developed countries. Int. J. Health Plann. Mgmt., 10: 113–128. doi: 10.1002/hpm.4740100205.

Howard, G. & Bartram, J. Domestic Water Quantity, Service Level and Health. World Health Organization, 2003. Retrieved on October 31, 2012 at http://cdrwww.who.int/water_sanitation_health/diseases/WSH03.02.pdf

United Nations Development Programme. The Human Poverty Index. Retrieved on October 31, 2012 at http://hdr.undp.org/en/statistics/indices/hpi/

The Sphere Project. Water supply standard 1: Access and water quantity. Retrieved on October 31, 2012 at http://www.spherehandbook.org/en/water-supply-standard-1-access-and-water-quantity/

Moser, C. O., & McIlwaine, C. (2006). Latin American urban violence as a development concern: towards a framework for violence reduction. World development, 34(1), 89-112.

Ghanizada. Security to be doubled for Afghan infrastructure projects. August 24, 2012. Khaama Press. Retrieved at http://www.khaama.com/security-to-be-doubled-for-afghan-infrastructure-projects/

U.S. Army, 2006. Counterinsurgency. Field Manual 3-24. Department of the Army, Washington, D.C. December, 2006.

Reuters. US Afghan aid could flop on counterinsurgency goals-watchdog. July 30, 2012. Retrieved at http://in.reuters.com/article/2012/07/30/usa-afghanistan-aid-idINL2E8IU61U20120730

Bonabeau, E. Agent-based modeling: Methods and techniques for simulating human systems. Proceedings of the National Academy of Sciences, vol. 99, no. Suppl. 3, pp. 7280-7287, 2002.

Cressie, N. Statistics for spatial data, revised edition. New York, NY.: John Wiley & Sons, Inc., 1993.

Macal C. & North, M. Tutorial on agent-based modeling and simulation. Journal of Simulation, vol. 4, pp. 151-162, 2010.

Transitional Islamic State of Afghanistan Ministry of Health (2004). Afghanistan National Hospital Survey. Available: http://www.msh.org/afghanistan/pdf/Afghanistan_National_Hospital_Survey.pdf

Jalalabad. (2012). In Encyclopædia Britannica. Retrieved from http://www.britannica.com/EBchecked/topic/299643/Jalalabad

Malleson, N. RepastCity - A demo virtual city. 2011. [Online]. Available: http://portal.ncess.ac.uk/wiki/site/mass/repastcity.html

Becu, N., Neef, A., Schreinemachers, P. & Sangkapitux, C. Participatory computer simulation to support collective decision-making: Potential and limits of stakeholder involvement, Land Use Policy, Volume 25, Issue 4, October 2008, Pages 498-509, ISSN 0264-8377, 10.1016/j.landusepol.2007.11.002.

Barreteau, O., Bousquet, F. & Attonaty, J.-M. (2001) Role-playing games for opening the black box of MAS: method and teachings of its application to Senegal River Valley irrigated systems. Journal of Artificial Societies and Social Simulation, vol. 4, no. 2.

Barreteau, O., Le Page, C. & D'Aquino, P (2003). Role-Playing Games, Models and Negotiation Processes. Journal of Artificial Societies and Social Simulation vol. 6, no. 2.

Appendix A. Design Points

Summary of design points in Section 5 (Results).

Location	Design	Replicates
1	Factorial	3
2	Factorial	3
3	Factorial	3
4	Factorial	3
5	Steepest Descent	3
6	Steepest Descent and CCD	3
7	Steepest Descent	3
8	Steepest Descent	3
9	Steepest Descent	3
10	CCD	3
11	CCD	3
12	CCD	3
13	CCD	3
14	CCD	3
15	CCD	3
16	CCD	3
17	CCD	3

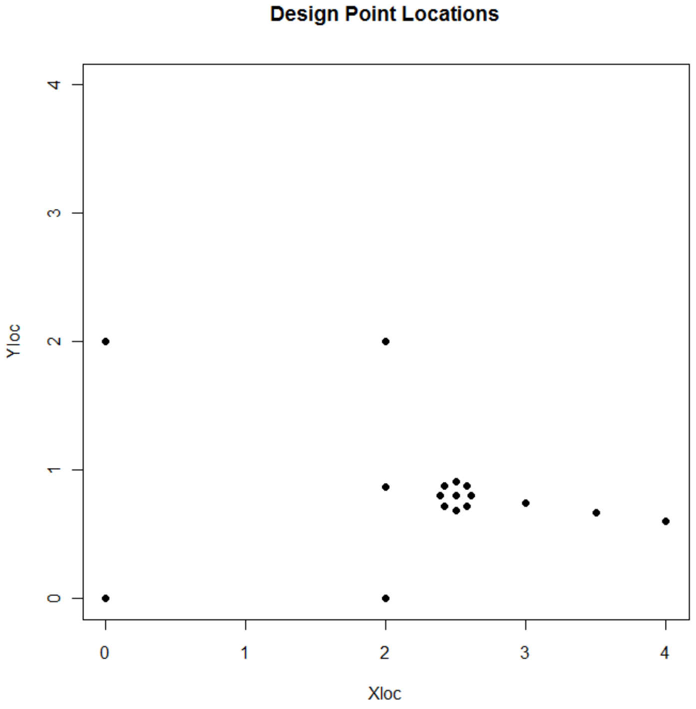

Figure 7 All seventeen of the design points used in Section 5

Appendix B: Raw Data

Raw data used in Section 5 (Results).

Design	Location	Xloc	Yloc	Out
Factorial Design Points	1	0	0	120
	1	0	0	144
	1	0	0	129
	2	2	0	108
	2	2	0	118
	2	2	0	143
	3	2	2	125
	3	2	2	123
	3	2	2	122
	4	0	2	139
	4	0	2	126
	4	0	2	134
Steepest Descent Design Points	5	2	0.87	108
	5	2	0.87	119
	5	2	0.87	124
	6	2.5	0.8	109
	6	2.5	0.8	122
	6	2.5	0.8	106
	7	3	0.74	110
	7	3	0.74	116
	7	3	0.74	121
	8	3.5	0.67	120
	8	3.5	0.67	128
	8	3.5	0.67	127
	9	4	0.6	140
	9	4	0.6	148
	9	4	0.6	142
CCD Design Points	6	2.5	0.8	109
	6	2.5	0.8	122
	6	2.5	0.8	106
	10	2.58	0.88	102
	10	2.58	0.88	106
	10	2.58	0.88	109
	11	2.42	0.88	120
	11	2.42	0.88	121
	11	2.42	0.88	129
	12	2.58	0.72	102
	12	2.58	0.72	99

12	2.58	0.72	106
13	2.42	0.72	113
13	2.42	0.72	110
13	2.42	0.72	112
14	2.613137	0.8	98
14	2.613137	0.8	98
14	2.613137	0.8	88
15	2.386863	0.8	99
15	2.386863	0.8	100
15	2.386863	0.8	93
16	2.5	0.913137	106
16	2.5	0.913137	103
16	2.5	0.913137	115
17	2.5	0.686863	122
17	2.5	0.686863	121
17	2.5	0.686863	130

Appendix C. Linear Model Details and Analysis

Factorial Model

```
lm(formula = data$Out ~ data$Xloc + data$Yloc, data = data)
Residuals:
    Min      1Q  Median      3Q     Max
-14.583  -5.083  -1.250   2.667  20.417

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) 131.4167     5.3664  24.489 1.51e-09 ***
data$Xloc    -4.4167     3.0983  -1.426    0.188
data$Yloc     0.5833     3.0983   0.188    0.855
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 10.73 on 9 degrees of freedom
Multiple R-squared: 0.1868,     Adjusted R-squared: 0.006101
F-statistic: 1.034 on 2 and 9 DF,  p-value: 0.3943
```

Steepest Descent

The averages of the three runs at each of the five steepest descent design points. The minimum (112.33 outpatients) occurs at design point 6.

Design Point	Average # of Outpatients
5	117
6	*112.3333*
7	115.6667
8	125
9	143.3333

CCD Model

```
lm(formula = Out ~ Xloc + Yloc + Xloc:Yloc, data = data)
Residuals:
      Min         1Q    Median         3Q       Max
-21.29149  -5.41221   0.09024   7.16050  19.71647

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)   -411.9      963.2  -0.428    0.673
Xloc           212.3      385.2   0.551    0.587
Yloc           801.1     1201.0   0.667    0.511
Xloc:Yloc     -325.5      480.3  -0.678    0.505

Residual standard error: 10.65 on 23 degrees of freedom
Multiple R-squared: 0.1421,     Adjusted R-squared: 0.03015
F-statistic: 1.269 on 3 and 23 DF,  p-value: 0.3083
```

Figure 8 shows a good Normal Q-Q plot and no outliers with significant leverage.

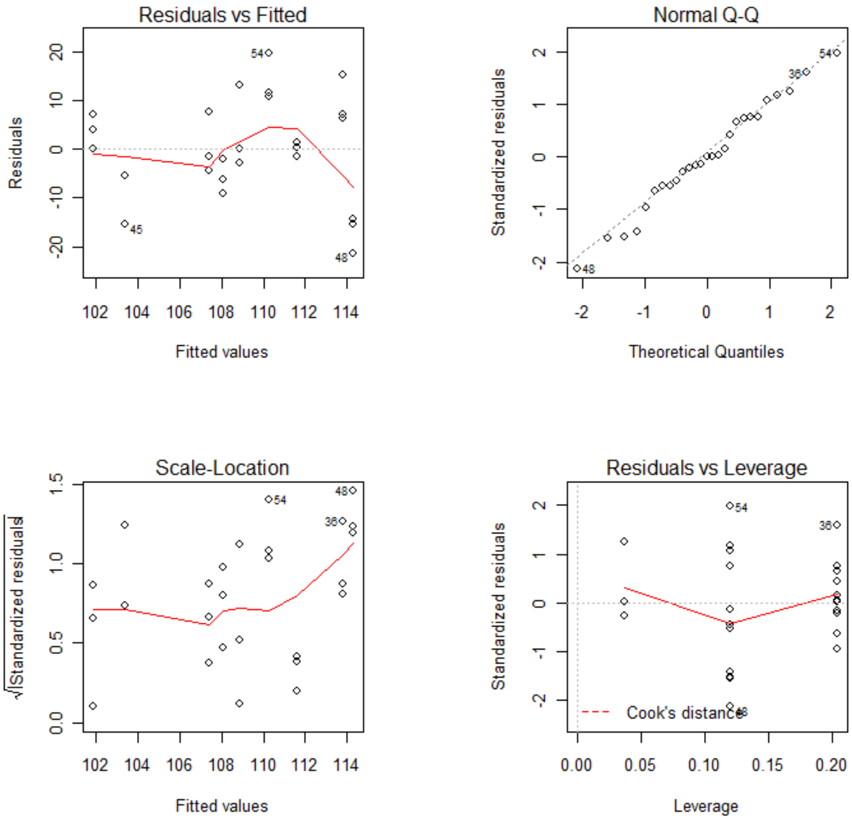

Figure 8 CCD model diagnostic plots.