

ARMY RESEARCH LABORATORY



Target Acquisition for Projectile Vision-Based Navigation

by Moshe Hamaoui

ARL-TR-6860

March 2014

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5066

ARL-TR-6860

March 2014

Target Acquisition for Projectile Vision-Based Navigation

Moshe Hamaoui

Weapons and Materials Research Directorate, ARL

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) March 2014		2. REPORT TYPE Final		3. DATES COVERED (From - To) June 2012–June 2013	
4. TITLE AND SUBTITLE Target Acquisition for Projectile Vision-Based Navigation				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Moshe Hamaoui				5d. PROJECT NUMBER AH80	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: RDRL-WML-F Aberdeen Proving Ground, MD 21005-5066				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-6860	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The application of computer vision to projectile navigation is a novel and untested proposition. As such, it is natural to wonder at this early stage how optical constraints may impact system performance. In this report, we develop a model for visual acquisition for a strap-down forward-facing imager with a constant field of view. Several prototypical scenarios are modeled, including stationary and moving target engagements. We further address the question of where to optimally place the initial ballistic impact point in relation to targets of uncertain position and heading. Finally, results of Monte Carlo simulations are presented to characterize the effects of target location error on visual acquisition.					
15. SUBJECT TERMS vision-based navigation, smart munitions, target acquisition, strap-down imager, machine vision					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 48	19a. NAME OF RESPONSIBLE PERSON Moshe Hamaoui
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code) 410-306-0968

Contents

List of Figures	iv
List of Tables	v
Acknowledgments	vi
1. Introduction	1
2. Projectile Model	2
2.1 Equations of Motion.....	4
2.2 Atmospheric Model.....	4
2.3 Initial Conditions.....	5
3. Target Model	7
3.1 Moving Target 1 (No Preferred Heading).....	7
3.2 Moving Target 2 (Preferred Heading).....	8
4. Where to Aim	10
4.1 No Preferred Heading.....	10
4.2 Preferred Heading.....	11
5. Simulation Studies	17
6. Discussion	18
7. Conclusion and Future Work	20
8. References	21
Appendix A. Simulation Results	23
Appendix B. Derivation of Ground Resolution for a Diffraction-Limited Pinhole Camera Model	35
Distribution List	40

List of Figures

Figure 1. Three critical flight stages for VBN: visual acquisition, feature recognition, and point of last maneuver.....	2
Figure 2. Geometry of visual acquisition.....	3
Figure 3. Initial conditions.....	5
Figure 4. Gamma distribution for target speed.	8
Figure 5. Target location distribution at time of projectile impact. Red arrow is drawn from reported position at time zero to mean target position at impact.	9
Figure 6. Probability of wraparound as the standard deviation of ϕ increases.....	10
Figure 7. Moving target distribution evolving in time.....	11
Figure 8. To ensure that all targets converge on the same ballistic impact point, it is necessary to offset the reported target position for each reported heading.	16
Figure 9. Representative results for Monte Carlo simulation of moving target 2 (with preferred heading).	16
Figure 10. Representative results for Monte Carlo simulation of moving target 2: middle 70th percentile bounds on α	17
Figure 11. Scene recognition is correlated with pixel resolution.....	18
Figure 12. Side-by-side results for visual acquisition (left) and target recognition (right).	19
Figure B-1. Differential object and image areas for pinhole camera model (left). Relation of field of view (FOV) to focal length and sensor dimensions (right).	36
Figure B-2. Resolution is proportional to the cube of the secant of the angular displacement from the optical axis (left). Equal areas exhibit pincushion distortion under transformation (right).	38
Figure B-3. Minimum resolvable target dimension for a specified FOV and slant range (left). Increased off axis resolution allows for smaller target recognition (right).	39

List of Tables

Table 1. Constraints of vision-based navigation.....	1
Table 2. Example 120-mm mortar physical data.	6
Table 3. Atmospheric data and physical constants.	6
Table 4. Summary of error parameters (moving target 1).	8
Table 5. Summary of error parameters (moving target 2).	9
Table 6. (Left) Target distributions after 40 s for various σ_ϕ . Arrow is directed from the target's initial position at t_0 to the spatial median after 40 s. (Right) Cumulative distribution functions for three central locations.	14
Table 7. Example imager specifications and assumptions.....	19

Acknowledgments

I would like to thank Frank Fresconi for his guidance and support in this research effort. In addition, I benefited greatly from conversations with Andrew Thompson on statistical modeling of moving targets. Many thanks, as well, to Tom Harkins for a very thorough technical review.

1. Introduction

Smart weapons technology has made great strides in recent years, leveraging an array of onboard sensors to deliver real-time guidance and navigation capability to gun-launched projectiles. However, traditional reliance on GPS to infer relative position of projectile with respect to target remains a significant system vulnerability and operational constraint. Taking inspiration from biological systems in the animal world, researchers at the U.S. Army Research Laboratory have begun to develop vision-based navigation (VBN) for precision munitions. This technology may one day provide the means to defeat fast-moving and highly maneuverable targets and to operate in GPS-denied environments. Although VBN has been implemented to varying degrees onboard robots, unmanned aerial vehicles, and even missiles, application to low-cost gun-launched projectiles will require the development of tightly coupled vision processing, guidance, and control algorithms, in addition to a highly maneuverable airframe and supporting optics. A crucial part of this effort is the development of a quantitative description of the physical constraints that bound optimal vision performance.

VBN is fundamentally predicated on the ability to observe and recognize visual cues from the surrounding environment. In the context of precision-guided munitions, these cues are typically few and far away. Furthermore, the projectile itself is hardly an ideal viewing platform. Onboard optics must withstand the violence of gun launch, placing severe limitations on mechanical swivel and refocusing systems. Once airborne, a forward-facing imager may be sky-stranded for the up leg of ballistic trajectories. Once the target comes into the imager's field of view (FOV), the projectile may yet be kilometers away. At this distance the target may not register as more than a few blurred pixels. Cloud cover, haze, aero-optical effects, atmospheric turbulence, and poor lighting can further obscure the target and complicate the task of image capture and recognition. As the projectile closes the intervening distance, the target continues to clarify. But whether effective action can be taken will depend on whether image capture and target recognition can be achieved while sufficient control authority yet exists.

The window of opportunity during which VBN is practically relevant is constrained by several factors, summarized in table 1.

Table 1. Constraints of vision-based navigation.

Situational Constraints	Device Constraints	Algorithmic Constraints
Projectile trajectory and orientation	Sensor resolution	Resolution threshold (i.e., pixels on target)
Slant range	Lens characteristics (e.g., FOV)	Contrast threshold
Target geometry and motion	Processor speed	Computational burden
Atmospheric conditions (visibility, turbulence, lighting) and aero-optical effects	—	—

These constraints naturally suggest three distinct milestones along the flight path, depicted schematically in figure 1:

1. Visual acquisition (VA)
2. Feature recognition (FR)
3. Point of last maneuver (PLM) (beyond which control authority has been relinquished)




Figure 1. Three critical flight stages for VBN: visual acquisition, feature recognition, and point of last maneuver.

These critical stages provide a useful framework in which to consider the performance of potential projectile VBN systems. Our goal, then, is to develop a set of models that describe the effects of situational, algorithmic, and device constraints on visual acquisition and feature recognition. These models can then be used to optimize design parameters, evaluate performance, and expose operational limitations.

This report will focus on the first of these critical points, visual acquisition. Several prototypical scenarios will be modeled, including stationary and moving target engagements.

2. Projectile Model

Visual acquisition refers, here, to the trajectory point at which the target first enters the camera's FOV. For a boresighted imager, the wider the FOV, the earlier visual acquisition can occur. However, maximizing the FOV comes at the expense of decreased image resolution and, hence, delayed feature recognition. At the opposite extreme, minimizing the FOV will increase

resolution (much like a telescope). This approach may be suitable for low quadrant elevation (QE) or flat-fire applications.* For indirect fire, however, this could delay visual acquisition well past the point of last maneuver. Instead, the optimal FOV must achieve a balance of sufficient resolution and early visual acquisition. To investigate this further, we first need a quantitative model of ballistic visual acquisition that incorporates realistic flight dynamics.

To simplify the mathematical description, we adopt the following assumptions:

- The projectile-mounted camera is forward facing, directed along the projectile's axis of symmetry.
- The projectile remains oriented along the instantaneous velocity vector, $\mathbf{v}$; there is no pitching or yawing motion.

Referring to figure 2, we can define the condition of visual acquisition as the instant when

$$\mathbf{v} \cdot \mathbf{s} = v s \cos\left(\frac{\theta_{FOV}}{2}\right) \quad (1)$$

where $\mathbf{v}$ is the velocity vector, $\mathbf{s}$ is the line-of-sight vector from projectile to target, and θ_{FOV} is the optical FOV. The optimal θ_{FOV} is not known a priori. Still, at any point along the trajectory, there will be a particular choice of θ_{FOV} that exactly satisfies the visual acquisition condition:

$$\theta_{FOV_{threshold}} = 2 \cos^{-1} \left[\frac{\mathbf{v} \cdot \mathbf{s}}{v s} \right] \quad (2)$$




Figure 2. Geometry of visual acquisition.

For convenience, we define a target angle, α , which describes the angle between the velocity and line-of-sight vectors.

$$\alpha \equiv \cos^{-1} \frac{\mathbf{v} \cdot \mathbf{s}}{v s} \quad (3)$$

* As a matter of practical consideration, a very narrow FOV could invite increased image jitter, pixel blur, and feature occlusion due to platform instability or other periodic motion.

To achieve visual acquisition at any time, t , the minimum θ_{FOV} is then

$$\theta_{FOV} = 2 \alpha[t] \quad (4)$$

Next, we must define $\alpha[t]$ in terms of QE, initial velocity, atmospheric conditions, and choice of airframe.

2.1 Equations of Motion

To model flight trajectories, we use a 3-degree-of-freedom point-mass model governed by the following set of coupled nonlinear ordinary differential equations (Fresconi et al., 2011):

$$\begin{aligned} \ddot{x} &= \frac{\pi \rho D^2 C_{x0} v}{8m} \dot{x} \\ \ddot{y} &= \frac{\pi \rho D^2 C_{x0} v}{8m} \dot{y} \\ \ddot{z} &= \frac{\pi \rho D^2 C_{x0} v}{8m} \dot{z} + g \end{aligned} \quad (5)$$

where the drag coefficient C_{x0} is generally a function of Mach number, M , specific to the choice of projectile. $M = v/c$. ρ and c (air density and speed of sound, respectively) generally vary with altitude.

2.2 Atmospheric Model

For projectile applications, we may safely assume the flight path is constrained within the troposphere, the bottommost portion of the Earth's atmosphere, extending to approximately 11 km above sea level. In this region, temperature is found to decrease linearly with altitude and can be empirically characterized by a constant lapse rate, L (in kelvins per meter). Following the International Standard Atmosphere model,^{*} we adopt the following assumptions:

- Hydrostatic equilibrium
- Perfect gas
- Gravitational acceleration constant with altitude
- Constant temperature lapse rate

Temperature and pressure can then be approximated with respect to ground pressure and temperature, T_0 and p_0 , (Portland State Aerospace Society, 2004):

$$T = T_0 - L z \quad (6)$$

^{*}Formulated by the International Civil Aviation Organization.

$$p = p_0 \left(1 - \frac{Lz}{T_0}\right)^{\frac{g\mu}{R}} \quad (7)$$

where p is ambient pressure, R is the universal gas constant, and μ is the molar mass of dry air. Having found T and p , we invoke the ideal gas law to solve for the air density, ρ :

$$\begin{aligned} \rho &= \frac{p\mu}{RT} \\ &= \frac{\mu p_0 \left(1 - \frac{Lz}{T_0}\right)^{\frac{g\mu}{R}}}{R(T_0 - Lz)} \end{aligned} \quad (8)$$

Finally, the speed of sound is given by:

$$\begin{aligned} c &= \sqrt{\gamma \frac{p}{\rho}} \\ &= \sqrt{\frac{\gamma R}{\mu} (T_0 - Lz)} \end{aligned} \quad (9)$$

where γ is the adiabatic index.

2.3 Initial Conditions

From figure 3, we can immediately write down the initial conditions in terms of the zenith and azimuth angles (equation 10):




Figure 3. Initial conditions.

$$\begin{aligned} x[0] &= 0 \\ \dot{x}[0] &= v_0 \sin[\theta_0] \cos[\phi_0] \\ y[0] &= 0 \\ \dot{y}[0] &= v_0 \sin[\theta_0] \sin[\phi_0] \\ z[0] &= 0 \\ \dot{z}[0] &= v_0 \cos[\theta_0] \end{aligned} \quad (10)$$

Note that $\dot{\phi} = 0$ (by assumption), so we may choose any constant ϕ_0 without loss of generality. We choose $\phi_0 \equiv 0$ so that $\phi[t] = 0$. The projectile-specific critical parameters are now m , v_0 , D , and θ_0 . Assuming no target location error (TLE), these four parameters will determine the trajectory and, in particular, $\alpha[t]$.

As an example, we select a 120-mm mortar. Tables 2 and 3 list some of the relevant physical parameters.

Table 2. Example 120-mm mortar physical data.

Parameter	Description	Value	Units (SI)
D	Diameter	0.120	m
m	Mass	13.68490	kg

Table 3. Atmospheric data and physical constants.

Parameter	Description	Value	Units
p_0	Pressure at sea level	101325	Pa
T_0	Temperature at sea level	288.15	K
L	Temperature lapse rate	0.0065	K/m
μ	Molar mass of dry air	0.0289644	kg/mol
γ	Ratio of heat capacities	7/5	—
R	Universal gas constant	8.31447	J/(K mol)
g	Acceleration due to gravity at sea level	9.80665	m/s ²

The system of differential equations (5) are solved numerically* with initial conditions (10) for various combinations of v_0 and θ_0 . To illustrate the response of α with respect to time and slant range, we choose three initial velocities and three QEs. v_0 values were 100, 200, and 300 m/s. QE values were 45° (800 mil), 65° (1156 mil), and 85° (1511 mil).

*This coupled set of nonlinear ordinary differential equations admits no analytic closed-form solution. See, for example, Taylor (2005, p 62).

3. Target Model

Perhaps the single greatest promise of VBN is the potential to engage moving targets whose relative trajectory is not absolutely known prior to firing.* The moving target scenario can take several forms, depending on the uncertainty of target location, speed, and heading. But while we could speak theoretically about Boolean combinations of these parameters, any bounded target distribution will require at least some information about probable location and speed. In this sense, there are only two distinct scenarios that we can investigate:

1. Probable location and speed (no reported heading information).
2. Probable location, speed, and heading.

To be sure, the source of the information (forward observer with laser range finder, plausible assumptions about mounted or unmounted infantry on rough terrain, etc.) is critical, as the probability density function describing the target location should correspond to the way in which the information was gathered. Still, for purposes of simulation, these differences are distinguishable only in the underlying probability distributions for location, speed, and heading. Our approach, then, shall be to investigate these two representative scenarios. The tool developed here should be adaptable to a range of operational situations with modification of the underlying probability distributions.

3.1 Moving Target 1 (No Preferred Heading)

In this first case, the initial target location and speed are known with some measure of certainty, while heading is entirely unspecified. Uncertainty in target position (like all measurements) will generally depend on the particular method of measurement employed by the observer and could result in higher fidelity along one orthogonal axis. For now, we assume a symmetric binormal distribution with the mean position at the last reported target location. Heading is assigned a uniform distribution (0° to 360°), while speed shall be defined by a gamma distribution[†] with a prescribed mean and standard deviation, as illustrated in figure 4. The target begins moving at the time of firing[‡] and thereafter maintains constant velocity. We are primarily interested in the earliest instant of visual acquisition, before corrective action has been taken. Effects of subsequent course-correcting maneuvers are therefore not considered.

*This would include cases of not only unspecified target position, but also of uncertain projectile position—hence the very important application to GPS-denied environments.

[†]Without experimental data, the decision to choose one distribution over another is more a matter of convenience than physical motivation. The gamma distribution, commonly parameterized by α and β , can be assigned an exact mean and standard deviation using $\alpha = \frac{\mu^2}{\sigma^2}$ and $\beta = \frac{\sigma^2}{\mu}$. Also, the support for this distribution, $x \in (0, \infty)$, correctly excludes nonphysical “negative” speeds.

[‡]This implicitly assumes reported information is still fresh at the time of firing. If this is not the case, it may be more appropriate to start the target moving at some prior time, yielding a proportional increase in TLE.




Figure 4. Gamma distribution for target speed.

The error parameters used in this simulation are summarized in table 4, and the results are collected in appendix A (section A.2).

Table 4. Summary of error parameters (moving target 1).

Symbol	Value	Notes
μ_{pos}	0 m	Mean displacement of initial target location from reported position, along each orthogonal axis in the XY (ground) plane
σ_{pos}	10 m	Standard deviation of target position along each coordinate in the XY plane
μ_{vel}	5 m/s	Mean target speed
σ_{vel}	2 m/s	Standard deviation of target speed

3.2 Moving Target 2 (Preferred Heading)

We now consider a scenario wherein the forward observer has reported not only speed and position, but also heading. Again, the reported information is necessarily uncertain because of measurement error. Figure 5 shows a representative trajectory and eventual target distribution (at the time of impact).




Figure 5. Target location distribution at time of projectile impact. Red arrow is drawn from reported position at time zero to mean target position at impact.

Error levels for position and speed are unchanged, but we have introduced additional parameters for the mean and standard deviation of the target heading, as shown in table 5.

Table 5. Summary of error parameters (moving target 2).

Symbol	Value	Notes
μ_{pos}	0 m	Mean displacement of initial target location from reported position, along each orthogonal axis in the XY (ground) plane
σ_{pos}	10 m	Standard deviation of target position along each coordinate in the XY plane
μ_{vel}	5 m/s	Mean of target speed
σ_{vel}	2 m/s	Standard deviation of target speed
μ_ϕ	0°, 45°, 90°, 135°, 180°	Mean azimuthal target velocity heading
σ_ϕ	15°	Standard deviation of azimuthal target velocity heading

The additional heading error is assumed to have a normal distribution. Technically, modeling a preferred heading with the unbounded normal distribution will allow for “wraparound” effects, such that the resulting standard deviation of the angular distribution will not exactly match the specified standard deviation. Practically, the contributions from the tail ends are completely negligible for any reasonable σ_ϕ . In particular, the probability of a wraparound sample point occurrence would be

$$\begin{aligned}
 P(\phi < \mu - \pi \vee \phi > \mu + \pi) &= 1 - \int_{\mu - \pi}^{\mu + \pi} \frac{e^{-\frac{(\phi - \mu)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} d\phi \\
 &= 1 - \operatorname{erf}\left(\frac{\pi}{\sqrt{2}\sigma}\right)
 \end{aligned} \tag{11}$$

Numerical evaluation shows that, for the assumed 15° standard deviation, the probability of wraparound is effectively zero (figure 6). This is not surprising, as the wraparound point (i.e., $\pm 180^\circ$) is 12 standard deviations from the mean.*




Figure 6. Probability of wraparound as the standard deviation of ϕ increases.

4. Where to Aim

Target location distribution (as a function of time) has been described relative to reported position. In order to calculate α , we must also fix the projectile's trajectory in relation to this point. In practical terms, the question is, Where should one aim in relation to the reported target position?

4.1 No Preferred Heading

Without any reported heading information, it may seem reasonable to shoot for the last known target position. However, the expanding “doughnut” distribution, depicted in figure 7, shows that a bull’s-eye shot would almost certainly miss the target. For a precisely aimed unguided ballistic trajectory, the last reported position of a moving target is probably a poor choice.[†]

*For applications with a much greater angular standard deviation, this wraparound effect could be avoided by judicious choice of an alternative *bounded* distribution.

[†]By “precisely aimed,” we mean that the expected miss distance (from intended impact point) is much less than the expected target displacement from the reported position. Of course, if the lethal effects radius is large enough to include a good part of the target distribution, then a dead-center miss may actually have the greatest kill probability.





Figure 7. Moving target distribution evolving in time.

In the context of guided munitions, however, a no-hit initial trajectory may be acceptable if this choice increases the probability of a successful follow-on corrective maneuver and ultimate target intercept. An optimal trajectory, seeking to maximize this probability, will generally depend on the control authority of the particular vehicle and the character of the target distribution. As a first-order approximation, we shall assume control authority is symmetric about the point of impact, and corrective action is proportional to expected miss distance, so that only the target distribution need be considered. It seems natural, then, to place the ballistic point of impact at the center of this distribution.* In subsequent targeting scenarios, it will be necessary to define this central point more carefully. For now, because we have a circular distribution with no preferred heading, this point is trivially located at the last reported target location.

4.2 Preferred Heading

As in the previous section, we must address the question of where to fix the initial ballistic impact point in relation to the reported target position. Again, we select the center of the target distribution, $\mathbf{p}_c$. However, as there is now a preferred heading, the target distribution is no longer

*The center of the annulus in figure 7 is clearly coincident with the centroid $(\bar{x}, \bar{y})$. It is not immediately apparent, however, that this point should also represent the spatial median (equation 12). Perhaps there exists some other point (or locus of points) that better minimizes the sum of distances. Consider the following figure:



By inspection, any point on segment $\overline{A_1A_2}$ is a spatial median of points $\{A_1, A_2\}$. By the same argument, if point F is included in all line segments $\overline{P_1P_2}$, then F is the unique spatial median of all points P_i . The radial symmetry of the target distribution allows us to pair off all points about the center, confirming that the geometric center is indeed coincident with the spatial median.

symmetric about the reported position, and we must more precisely define the distribution center. Notionally, the motivation for using a “center point” is that, on average, it’s “closest” to all other points and would therefore tend to capture a maximum of the probability density distribution (of target location) within the control shadow of the projectile. In this sense, what we need to find is the coordinate from which the sum of Euclidean distances to all other distribution points is minimized, the so-called *spatial median*.^{*} Specifically,

$$\mu_{sm}(X) = \underset{\mu \in \mathbb{R}^p}{\operatorname{argmin}} \sum_{i=1}^n \|x_i - \mu\| \quad (12)$$

where X is a set of points in $\mathbb{R}^p$, and $\mu_{sm}(X)$ is the spatial median of the set. For a target constrained to the XY plane (bivariate case, no flight capability, hills, or valleys), $p = 2$.

Historically, the analogous problem was first posed by Fermat for three noncollinear points and exactly solved by Torricelli in the mid 17th century (Hale and Moberg, 2003). The more general problem of n points in two or more dimensions was formulated by Weber in 1909, with reference to the problem of optimal facility location (Hale and Moberg, 2003), but no closed-form solution exists, and algorithmic approximations are computationally intensive (Bose et al., 2003). From a practical standpoint, the use of the spatial median to calculate and update the optimal ballistic impact point may not be well suited for real-time applications.[†]

A popular alternative distribution center is given by the centroid, which minimizes the sum of squared Euclidean distances between itself and each point in a set. While the centroid is the physically correct choice for many applications, it is not generally coincident with the spatial median and would therefore serve only as an approximation of the desired central tendency. Still, the centroid can be calculated in linear time for a discrete set of points and may, as in this case, admit closed-form solutions if the probability density function is known. In the present scenario, the target motion may be described by the vector equation:

$$\mathbf{p}_{target} = \mathbf{p}_{initial} + \mathbf{v} t \quad (13)$$

Equivalently, in component form, we have:

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x_i \\ y_i \end{pmatrix} + \begin{pmatrix} \cos \phi \\ \sin \phi \end{pmatrix} vt \quad (14)$$

where v is the scalar target velocity, and ϕ is the (azimuthal) target heading.

^{*}Also referred to as the *geometric median* or *Fermat-Weber point*.

[†]Much research has gone into speeding up approximation algorithms for the spatial median of an arbitrary set of points. See, for example, Bose et al. (2003). In the present context, however, it seems reasonable to represent the central target location in the form $\mathbf{p}_{central} = \mathbf{p}_{initial} + \beta \mathbf{\mu}_v t$, where $\mathbf{\mu}_v$ is the mean (reported) velocity, and α is a constant, which depends only on σ_ϕ and the median type (e.g., centroid, L1, L2). Indeed, for the centroid descriptor we shall derive a closed form for $\beta[\sigma_\phi]$. For other descriptors, this table of coefficients could be compiled numerically.

In general, for two independent random variables $\{P, Q\}$ and constants $\{a, b\}$, the expectations will follow (Zwillinger, 2003):

$$\begin{aligned} E[aP + bQ] &= aE[P] + bE[Q] \\ E[PQ] &= E[P]E[Q] \end{aligned} \quad (15)$$

Then for target position $\mathbf{p}_{target}$, we have:

$$\begin{aligned} E \left[\begin{pmatrix} X \\ Y \end{pmatrix} \right] &= E \left[\begin{pmatrix} X_i \\ Y_i \end{pmatrix} \right] + E \left[\begin{pmatrix} \cos \Phi \\ \sin \Phi \end{pmatrix} V \right] t \\ &= E \left[\begin{pmatrix} X_i \\ Y_i \end{pmatrix} \right] + E \left[\begin{pmatrix} \cos \Phi \\ \sin \Phi \end{pmatrix} \right] * E[V]t \end{aligned} \quad (16)$$

Now, $E \left[\begin{pmatrix} X_i \\ Y_i \end{pmatrix} \right]$ and $E[V]$ correspond to the reported target position and speed, respectively.

Finally, $E \left[\begin{pmatrix} \cos \Phi \\ \sin \Phi \end{pmatrix} \right]$ can be calculated for each component. By construction, Φ draws from a normal distribution. The expectation (mean) along each dimension is then*

$$\begin{aligned} E(\cos \Phi) &= \int_{-\infty}^{\infty} \cos \phi * \frac{e^{-\frac{(\mu-\phi)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} d\phi \\ &= e^{-\frac{\sigma^2}{2}} \cos \mu \end{aligned} \quad (17)$$

$$\begin{aligned} E(\sin \Phi) &= \int_{-\infty}^{\infty} \sin \phi * \frac{e^{-\frac{(\mu-\phi)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} d\phi \\ &= e^{-\frac{\sigma^2}{2}} \sin \mu \end{aligned} \quad (18)$$

The centroid of the target distribution is composed of the mean along each dimension:

$$\mathbf{p}_{centroid} = \mathbf{p}_{reported} + \begin{pmatrix} e^{-\frac{\sigma_\phi^2}{2}} \cos \mu_\phi \\ e^{-\frac{\sigma_\phi^2}{2}} \sin \mu_\phi \end{pmatrix} v_{reported} t \quad (19)$$

An alternative central location can be defined by the coordinate-wise median along each dimension. This point exactly solves the Weber problem under the L_1 (Manhattan) metric (Bose et al., 2003). Given the inherent symmetry of the target distribution and the uniqueness of the $n \geq 2$ dimensional spatial median (Milasevic and Ducharme, 1987), the spatial median must lie along the axis of symmetry (i.e., the reported target heading vector). Furthermore, as the 2-D target distribution collapses toward the axis of symmetry, the spatial median must exactly approach the sample median along the 1-D downrange axis.[†] These observations suggest that for

*The ϕ subscripts are suppressed for readability.

[†]The spatial median of any even number of arbitrarily placed points along one dimension is not unique, but the locus of points satisfying this condition must include the sample median. For an odd number of points, the solution is unique and coincident with the median.

small σ_ϕ , the location of the spatial median may be approximated by the coordinate-wise median. Indeed, table 6 shows that for small σ_ϕ , the coordinate-wise median almost exactly overlays the spatial median and appears (even for larger σ_ϕ) to consistently provide a better fit than does the centroid.

Table 6. (Left) Target distributions after 40 s for various σ_ϕ . Arrow is directed from the target's initial position at t_0 to the spatial median after 40 s. (Right) Cumulative distribution functions for three central locations.




Table 6. (Left) Target distributions after 40 s for various σ_ϕ . Arrow is directed from the target's initial position at t_0 to the spatial median after 40 s. (Right) Cumulative distribution functions for three central locations (continued).




Table 6 compares the cumulative distribution function of miss distances for centroid, coordinate-wise median, and spatial median. Although the cumulative distribution function is sensitive to σ_ϕ , the particular choice for central estimate does not appear to have an appreciable effect for small σ_ϕ . For the present application, the spatial median provides only a small advantage over the other approximations just described. For computational convenience, we shall therefore employ the centroid formulation (equation 19) as our working central location estimate.

For each trajectory, characterized by $\{v_0, QE\}$, we would like to investigate visual acquisition performance as a function of reported target heading, μ_ϕ . From equation 19 it is clear that the displacement vector $\mathbf{p}_c[t_{\text{impact}}] - \mathbf{p}_c[t_0]$ will depend on μ_ϕ . In order to ensure that $\mathbf{p}_c[t_{\text{impact}}]$ remains identical across all headings, we must appropriately offset the reported target position.* From a modeling standpoint, this is accomplished by defining $\mathbf{p}_{\text{rep}}$ with respect to the eventual (fixed) ballistic impact point.

$$\mathbf{p}_{\text{rep}} = \mathbf{p}_{\text{ballistic}} - \begin{pmatrix} e^{-\frac{\sigma_\phi^2}{2}} \cos \mu_\phi \\ e^{-\frac{\sigma_\phi^2}{2}} \sin \mu_\phi \end{pmatrix} v_{\text{rep}} t_{\text{impact}} \quad (20)$$

In effect, the reported position must vary so that all targets converge on the same point (figure 8).

* Another way to maintain experimental uniformity would be to “report” all targets at the same initial position (with variable headings) and then have them propagate toward separate ballistic impact points. But this implies changing also the projectiles’ initial conditions $\{v_0, QE\}$ to assure an intercept.




Figure 8. To ensure that all targets converge on the same ballistic impact point, it is necessary to offset the reported target position for each reported heading.

It should be emphasized, therefore, that even without any TLE, the relative orientation of projectile and target (measured in terms of the angle α) will depend on target heading. In particular, because we have aligned the x axis along the downrange axis, a target heading of $\mu_\phi = 0^\circ$ will always yield the largest α across the trajectory, as illustrated in figure 9. Analogous results for all selected combinations of $\{v_0, QE, \mu_\phi\}$ are included in appendix A (section A.3.1).




Figure 9. Representative results for Monte Carlo simulation of moving target 2 (with preferred heading).

From a design and requirements standpoint, it is useful to capture the effects of TLE in the context of an idealized α history. Figure 10 shows performance by percentile as well as the no-TLE estimates for a representative case. Analogous results for all selected combinations of $\{v_0, QE, \mu_\phi\}$ are included in appendix A (section A.3.2).




Figure 10. Representative results for Monte Carlo simulation of moving target 2: middle 70th percentile bounds on α .

5. Simulation Studies

We consider a matrix of nine projectile trajectories, representing combinations of three possible QE angles and three initial velocities. Three scenarios are simulated:

- Stationary target without TLE
- Moving target without preferred heading (with and without TLE)
- Moving target with preferred heading (with and without TLE)

In all cases, projectile-state is assumed to be error-free. The simulation code has been written in *Mathematica* and is easily adapted to other projectile airframes, initial conditions, and target statistics.

Plots displaying results of Monte Carlo simulations are found in appendix A.

6. Discussion

The application of VBN to projectile navigation is a novel and untested proposition. As such, it is natural to wonder at this early stage whether optical considerations will prove to be a minor engineering detail or a central research challenge. On its own, the present analysis does not fully address this question. Overall performance will depend not only on visual acquisition, but also on target recognition (among other factors). Still, because sight must logically precede detection, the characterization of visual acquisition in the context of projectile and target dynamics places a definite and predictable limit on earliest possible detection. In the presence of TLE, especially for a fast-moving target, these geometric and statistical considerations also reveal a “latest possible detection.”

While a complete discussion of the physical constraints of target recognition is beyond the scope of this report, it is instructive to consider the simple case of a diffraction-limited pinhole camera, neglecting atmospheric and aero-optical effects. We assume a resolution threshold of 32×32 pixels on-target is needed for recognition.* The idea that scene recognition is directly correlated with pixels (or line pairs) on-target seems intuitive, as illustrated in figure 11.



Figure 11. Scene recognition is correlated with pixel resolution.

Device specifications and other assumptions are listed in table 7. The resulting minimal dimension for a discernible target is plotted with respect to slant range for several FOVs[†] (figure 12).

*This threshold was found empirically for human scene recognition (Torrvalba, 2009). Other recognition thresholds can be found in the literature—most notably, the Johnson criteria (Johnson, 1958), which requires six line pairs (or 12 pixels) along the minimum target dimension for a 50% probability of recognition. See Kopeika (1998) for a more complete discussion. The actual resolution threshold for *machine* vision will depend on the particular vision algorithms used and should be measured experimentally for a given system. However, the consensus appears to be that machine vision will not outperform human recognition in the foreseeable future. Our use of human threshold levels should therefore be viewed as a best-possible expectation for machine vision.

[†]A derivation of this resolution model is included in appendix B. The diffraction limit is given by the Rayleigh resolution criterion: two point sources of equal intensity can just be resolved with diffraction-limited optics if separated by a view angle of $\sin^{-1}(1.22 \lambda/D)$, where λ is the wavelength, and D is the aperture diameter (Woan, 2003).

Table 7. Example imager specifications and assumptions.

Imager Specifications	
Resolution	2592 × 1944
Diagonal field of view	50°
Aperture	3 mm
Assumptions	
λ	550 nm (green)
Threshold pixels on target	32 × 32




Figure 12. Side-by-side results for visual acquisition (left) and target recognition (right).

Inasmuch as the effective VBN window depends critically on initial conditions, especially QE, this particular instance shows the target does not come into view until a range of about 430 m. To be recognized at this distance, the target must have a minimum dimension of about 4 m. The last opportunity for recognition comes at about 250 m, by which point the minimum target dimension has dropped to 2.5 m. As such, these numbers might prove feasible; however, this assumes onboard processors can handle the 5 Mpix throughput in real time.* In fact, this turns out to be a dangerous assumption, as feature extraction algorithms typically require considerable computational overhead.

For comparison, Svab et al. (2009) report an implementation of the Speeded-up Robust Features algorithm on a Field Programmable Gate Array at high-definition (1024 × 768 pixels) resolution achieving 10 fps (frames per second). At this resolution, the minimum target dimension increases to approximately 10 m to achieve recognition at 430 m.

*To meet guidance and control requirements, a tactical VBN solution would need to operate at approximately 10–20 fps (Malley, 2013).

The preceding example serves to illustrate some of the performance trade-offs and design challenges that must be addressed to enable projectile VBN. It seems likely that a successful implementation of projectile VBN will require larger standoff visual acquisition against smaller targets. Several schemes have been considered toward this end,^{*} including the use of “view-steering,” which we briefly explore in the next section.

7. Conclusion and Future Work

In this study, we have developed a model for visual acquisition for a strap-down forward-facing imager with a fixed FOV. We hope that this analysis, together with other models being developed for machine recognition, can be used to optimize design parameters, evaluate performance, and expose operational limitations of vision-based projectile navigation systems. A subsequent analysis will investigate the physical constraints on target recognition arising from sensor and lens characteristics, as well as atmospheric effects and lighting conditions.

It is clear, however, that there must always be a fundamental trade-off between FOV (and hence visual acquisition) and resolution (and hence recognition). To be sure, larger pixel arrays can increase resolution somewhat, but the increased computational burden can be crippling for real-time operation on a low-cost mobile platform, as mentioned previously. This constraint derives from the necessity to search the full FOV to initially find the target. In fact, it may be possible to leverage onboard sensor information (from inertial measurement units, magnetometers, etc.) and minimal target information to approximate relative target position in body coordinates and further narrow the search region. Coupled with a view-steering mechanism (mechanical or digital), this could provide a powerful VBN technique worth further consideration.

^{*}One notable approach employs scaled templates. A single high-resolution template is continuously updated (resampled and cropped) to correspond to the smallest recognizable target. Thus, for example, the algorithm might first search for a 32×32 -pixel city block, then a 32×32 -pixel house, and so on.


8. References

- Bose, P.; Maheshwari, A.; Morin, P. Fast Approximations for Sums of Distances, Clustering and the Fermat–Weber Problem. *Computational Geometry* **2003**, 24 (3), 135–146.
- Fresconi, F.; Cooper, G.; Costello, M. Practical Assessment of Real-Time Impact Point Estimators for Smart Weapons. *Journal of Aerospace Engineering* **2011**, 24 (1), 1–11.
- Gustavsson, B. Camera Calibration Toolbox for Matlab; 2010. http://www.vision.caltech.edu/bouguetj/calib_doc/ (accessed 6 May 2013).
- Hale, T. S.; Moberg, C. R. Location Science Research: A Review. *Annals of Operations Research* **2003**, 123, 21–35.
- Johnson, J. Analysis of Image Forming Systems. Image Intensifier Symposium. Warfare Vision Branch, Electrical Engineering Department, U.S. Army Engineer Research and Development Laboratories Corps of Engineers, Ft. Belvoir, VA, 1958; pp 249–276.
- Kopeika, N. S. *A System Engineering Approach to Imaging*; SPIE: Bellingham, Washington, 1998.
- Malley, J. M. U.S. Army Research Laboratory, Aberdeen Proving Ground, MD. Private communication, December 2013.
- Milasevic, P.; Ducharme, G. R. Uniqueness of the Spatial Median. *The Annals of Statistics* **1987**, 15 (3), 1332–1333.
- Portland State Aerospace Society Web site. A Quick Derivation Relating Altitude to Air Pressure, 2004 December 22. <http://psas.pdx.edu/RocketScience/> (accessed 30 July 2012).
- Svab, J.; Krajnik, T.; Faigl, J.; Preucil, L. FPGA Based Speeded Up Robust Features. In *Technologies for Practical Robot Applications, 2009. TePRA 2009. IEEE International Conference on IEEE International Conference*, November 2009; pp 35–41.
- Taylor, J. R. *Classical Mechanics*; University Science Books: Sausalito, CA, 2005.
- Torralba, A. How Many Pixels Make an Image. *Visual Neuroscience* **2009**, 26 (1), 123–131.
- Woan, G. *The Cambridge Handbook of Physics Formulas*; Cambridge University Press: Cambridge, UK, 2003.
- Zwillinger, D. *CRC Standard Mathematical Tables and Formulae*; Chapman & Hall/CRC Press: Boca Raton, FL, 2003.

INTENTIONALLY LEFT BLANK.


Appendix A. Simulation Results

A.1 Stationary Target (no TLE)



A.2 Moving Target 1 (no preferred heading)

Target location error is introduced with respect to reported position and speed. In this scenario, target heading is unknown and assumed to be uniformly distributed. The notation is defined in table 4 of the main text.





A.3 Moving Target 2 (with preferred heading)


In this scenario, nominal target heading is reported. Target location error is introduced with respect to reported position, speed and heading. Notation is defined in table 5 of main text.


A.3.1 Median Response

The left-hand column shows results without TLE. The right-hand column shows the results for the corresponding Monte Carlo simulations with TLE. Only the median response is shown.












A.3.2 Percentile Distribution


Performance of the 15th, 50th (median), and 85th percentiles are plotted (orange, blue, and red, respectively) for each combination of target heading, initial projectile velocity, and QE . For comparison, the corresponding “no-error” α histories are superimposed (dashed line).











INTENTIONALLY LEFT BLANK.

Appendix B. Derivation of Ground Resolution for a Diffraction-Limited Pinhole Camera Model

We wish to define the resolution as the differential ratio of image pixels to object area dp/dA_o as a function of slant range and angular position (figure B-1). The object area is taken to be the target profile as seen from the imager.




Figure B-1. Differential object and image areas for pinhole camera model (left). Relation of field of view (FOV) to focal length and sensor dimensions (right).

We assume a pinhole camera model. The 3-D object coordinates $\{x, y, z\}$ are thus mapped to the 2-D image coordinates $\{x_i, y_i\}$ by the projection transformation

$$\begin{pmatrix} x_i \\ y_i \end{pmatrix} = -\frac{f}{z} \begin{pmatrix} x \\ y \end{pmatrix} \quad (\text{B-1})$$

where f is the focal length. Expressing $\{x, y\}$ in spherical coordinates yields

$$\begin{pmatrix} x_i \\ y_i \end{pmatrix} = -f \begin{pmatrix} \cos \phi \tan \theta \\ \sin \phi \tan \theta \end{pmatrix} \quad (\text{B-2})$$

The differential image area is then (making use of the Jacobian of the transformation)

$$\begin{aligned} dA_i &= \left| \frac{\partial(x, y)}{\partial(\theta, \phi)} \right| d\theta d\phi \\ &= \left| \det \begin{pmatrix} -f \cos \phi \sec^2 \theta & f \sin \phi \tan \theta \\ -f \sec^2 \theta \sin \phi & -f \cos \phi \tan \theta \end{pmatrix} \right| d\theta d\phi \\ &= f^2 |\sec^2 \theta \tan \theta| d\theta d\phi \end{aligned} \quad (\text{B-3})$$

The differential object area dA_o for a spherical shell of constant radius r is $r^2 \sin \theta d\theta d\phi$. The resolution can now be expressed.

$$\begin{aligned}
Res &= \frac{dp}{dA_o} \\
&= \left(\frac{dp}{dA_i} \right) \frac{dA_i}{dA_o} \\
&= \frac{p_h p_w}{hw} \frac{f^2 \sec^2 \theta \tan \theta d\theta d\phi}{r^2 \sin \theta d\theta d\phi} \\
&= \frac{p_h p_w}{hw} \frac{f^2 \sec^3 \theta}{r^2}
\end{aligned} \tag{B-4}$$

where $\{p_h, p_w\}$ are the height and width of the sensor array in pixels, respectively. Referring to figure B-1, we can express f as a function of FOV:

$$f = \frac{1}{2} \frac{\sqrt{h^2 + w^2}}{\tan(FOV/2)} \tag{B-5}$$

In terms of the FOV, the resolution is then

$$Res = \frac{p_h p_w (h^2 + w^2)}{4hwr^2} \frac{\sec^3 \theta}{\tan^2 \left(\frac{FOV}{2} \right)} \tag{B-6}$$

The resolution is not uniform over the imager but rather increases as $\sec^3 \theta$. Therefore, as the FOV increases, we can expect more pronounced border distortion* (figure B-2). With proper calibration, this effect, as well as other optical distortions causing radial and tangential image deformations, can be largely rectified in software.†

*For the pinhole camera model under consideration, this is characterized by so-called pincushion distortion, depicted in figure B-2. Lenses may exhibit various forms of radial distortion. A notable example is the fish-eye lens, which is characterized by decreasing border resolution.

†Several such camera-calibration and image post-processing packages exist, including open source codes, such as Gustavsson (2010). Proper calibration is particularly important for this application given the real-time and low-cost mission requirements.




Figure B-2. Resolution is proportional to the cube of the secant of the angular displacement from the optical axis (left). Equal areas exhibit pincushion distortion under transformation (right).

For a target profile with area d^2 , the number of pixels on the target is $(d^2 \times Res)$. Requiring p_{min}^2 pixels on the target and solving for d , we can express the minimum target dimension as a function of slant range:

$$d_{min} = \frac{2 \sqrt{h w} p_{min} \tan \frac{FOV}{2} (\cos \theta)^{3/2} r}{\sqrt{p_h p_w (h^2 + w^2)}} \quad (B-7)$$

For a square focal array with $p_w \times p_w$ pixels, equation B-7 reduces to

$$d_{min} = \frac{\sqrt{2} p_{min} \tan \frac{FOV}{2} (\cos \theta)^{3/2} r}{p_w}$$

As an example, consider an imager with a 1024- $\times$ 1024-pixel array. The dimension of the smallest resolvable target is plotted as a function of slant range for selected FOVs (figure B-3). Although increased off-axis resolution theoretically allows for smaller target recognition, these gains are somewhat offset by increased radial aberration.




Figure B-3. Minimum resolvable target dimension for a specified FOV and slant range (left). Increased off axis resolution allows for smaller target recognition (right).

NO. OF COPIES	ORGANIZATION
1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
1 (PDF)	DIRECTOR US ARMY RESEARCH LAB IMAL HRA
1 (PDF)	DIRECTOR US ARMY RESEARCH LAB RDRL CIO LL
1 (PDF)	GOVT PRINTG OFC A MALHOTRA
1 (PDF)	DIR USARL RDRL WML F M HAMAOU