doi: 10.82124/CEAS-GNC-2026-046 CEAS-GNC-2026-046

EuroGNC Navigation Guidance Control CEAS Madrid, Spain May 5th-7th uc3m Universidad Carlos III de Madrid AIAA 2026

Optimizing Proportional Navigation for Maximum Terminal Velocity and Real-Time Intercept Estimation

Hasan Uçar Guidance, Navigation and Control Engineer, Roketsan Inc., Strategic Air Defense Systems Design Department, Ankara, Türkiye. hasan.ucar@roketsan.com.tr

Murat Önen Guidance, Navigation and Control Engineer, Roketsan Inc., Strategic Air Defense Systems Design Department, Ankara, Türkiye. murat.onen@roketsan.com.tr
Middle East Technical University, Ankara, Türkiye. murat.onen@metu.edu.tr

İbrahim Emre Demir Guidance, Navigation and Control Engineer, Roketsan Inc., Strategic Air Defense Systems Design Department, Ankara, Türkiye. emre.demir.1@roketsan.com.tr
Middle East Technical University, Ankara, Türkiye. demir.emre_01@metu.edu.tr

ABSTRACT

Missile guidance accuracy and terminal performance are critical in engagements against maneuvering or time-sensitive targets. Achieving both high terminal impact velocity and reduced time of flight is essential to enhance terminal homing effectiveness. This work proposes a two-stage midcourse guidance strategy: (i) offline optimization of Proportional Navigation Guidance (PNG) parameters to maximize terminal velocity across varying engagement geometries, and (ii) real-time intercept-point estimation combined with lookup-table retrieval of near-optimal PNG gains to adapt dynamically during flight. The proposed method is evaluated by comparing estimated and actual impact conditions, demonstrating improved response time, higher terminal velocity, and robust performance relative to conventional PNG approaches.

Keywords: Proportional Navigation Guidance (PNG), Intercept Point Estimation, Terminal Velocity Maximization, Energy-Efficient Engagement, IPM-PNG

Nomenclature

p_m, p_t	=	missile and target position vectors
R, $\dot{R}$	=	relative range and its rate
λ, $\dot{\lambda}$	=	line-of-sight (LOS) angle and LOS rate
V, V_T, V_c	=	missile speed, target speed, and closing speed
N	=	PNG navigation constant (ratio)
a_c	=	commanded lateral acceleration
V_{impact}	=	terminal speed at intercept
d_{miss}	=	miss distance (constraint)
T	=	time of flight (TOF)
J	=	cost-function value
w_V, w_T	=	velocity and TOF cost weights ($w_V + w_T = 1$)

ε	=	miss-distance tolerance
$\mathcal{T}$	=	grid of candidate intercept states (x, z)
$\mathcal{N}$	=	sampled set of navigation constants
$\mathcal{F}$	=	LUT mapping $(x, z) \mapsto N_{\text{opt}}(x, z)$
ΔN	=	navigation-constant increment
N_{optimal}	=	optimal N at a grid point
$x,\ z$	=	downrange (horizontal distance from launch point) and altitude coordinates
(x^*, z^*)	=	estimated intercept coordinates
$i,\ k$	=	grid index and iteration counter
d_i	=	downrange grid value
t_i	=	TOF associated with d_i
D_i	=	computed intercept downrange
Δ_i	=	downrange mismatch $d_i - D_i$

1 Introduction

In modern air defense systems, the effectiveness of a surface-to-air missile (SAM) largely depends on its ability to engage high-speed aerial targets with precision and minimal reaction time. A critical component of this engagement is the midcourse guidance phase, which must guide the missile into the terminal phase with maximum velocity and within minimal time to ensure a successful intercept. Achieving this requires a careful balance between minimizing induced drag caused by unnecessary accelerations and time it takes to follow coarser paths.

In practice, optimal trajectories tend to deviate from simple linear paths, often following parabolic arcs that exploit variations in air density and pressure to minimize drag and maximize energy efficiency. However, due to the complexity and non-linearity of the system, closed-form solutions for velocity and time-of-flight optimization are generally not solvable. Conventional midcourse guidance strategies often rely on simplified models and assumptions that do not fully account for the non-linearities in aerodynamic behavior and atmospheric properties that vary with altitude. To address this, we propose a simulation-based optimization approach in which the missile's time-of-flight and final Mach number are precomputed over a set of predefined points in a 2D engagement domain. This allows for the offline construction of an optimal trajectory map, enabling real-time guidance to any point in the domain with known performance characteristics.

Moreover, by estimating the target's velocity and extrapolating its future position based on the time-of-flight data, the system can compute a predicted interception point, even when the target moves in an arbitrary direction. This methodology effectively reduces any engagement scenario to a 2D optimization problem, simplifying the guidance logic while maintaining high accuracy and performance.

This paper presents the formulation, simulation framework, and validation of this midcourse guidance strategy, demonstrating its effectiveness across a range of engagement scenarios with maneuvering targets and complex atmospheric conditions.

1.1 Problem Statement

Proportional Navigation Guidance (PNG) commands lateral acceleration to drive the line-of-sight (LOS) rate toward zero, guiding the missile to the target—moving or stationary—under ideal kinematics. In such ideal conditions (constant speeds, no actuator/propulsion limits, no atmosphere), PNG achieves negligible miss distance. In realistic flight, however, altitude-dependent aerodynamics, thrust variation, actuator limits, and guidance/avionics latencies make "zero-miss with minimum effort" misaligned with the operational objectives of high terminal speed and short time of flight (TOF).

Large, aggressive turns shorten TOF and suppress miss distance but increase drag and energy loss; shallow, energy-conserving trajectories preserve terminal speed at the cost of longer TOF. The navigation constant N therefore encodes a fundamental trade-off that is state- and time-dependent in practice. Because the relevant mappings—between altitude, Mach number, angle of attack, and drag/thrust—are highly nonlinear, linearized or closed-form treatments are unreliable for selecting N across the engagement envelope.

The problem addressed in this work is to determine, for each candidate intercept state, a midcourse PNG gain that preserves accuracy while maximizing terminal impact speed and controlling TOF, without resorting to heavy online optimization. We pursue a simulation-based approach: offline evaluation over a discretized engagement domain to capture the nonlinear physics, followed by real-time retrieval (with interpolation) of the near-optimal gain associated with the current intercept estimate. This frame of reference targets the 'sweet spot' between exploiting higher altitude, lower density flight for energy retention and limiting the commanded accelerations that would otherwise dissipate kinetic energy, thus extending effective range and improving intercept performance.

1.2 Related Literature

Proportional Navigation Guidance (PNG) drives the modern story of missile guidance, using a strategy rooted in a fundamentally simple principle: steering the line-of-sight (LOS) rate to zero. This fascinating concept revolutionized the field, and by the mid-1980s, analyses—most notably Guelman's groundbreaking work in 1984—had illuminated not only why PNG excels in idealized scenarios featuring minimal miss distances and computational simplicity but also foreshadowed the secrets hidden in more complex, real-world conditions influenced by atmospheric factors and actuator limitations [1]. This foundational chapter established PNG's reputation as robust and practical, yet it highlighted its indifference to energy optimization and the altitude-dependent physics that determine terminal performance.

Entering the 1990s, the paradigm evolved from a singular concentration on the 'hit' to a more strategic consideration of 'how to get there'. Kee and his collaborators (1998) redefined midcourse guidance as a near-time-optimal challenge, revealing that meticulously crafted trajectories could surpass the effectiveness of fixed-gain PNG while still being feasible in real-time applications [2]. This transformative period brought to light the operational constraints—actuator bandwidth, saturation limits, and envelope restrictions—that complicate the pursuit of direct optimal control onboard. By 2010, the emphasis transitioned to system-oriented insights, underscoring the synchronization of maneuvers, terminal constraints, and practical feasibility (as captured in Palumbo's 2010 work), thus reinforcing the necessity for midcourse strategies that harmonize with both hardware limitations and the ticking clock, rather than focusing exclusively on kinematics [3].

Simultaneously, the 2000s witnessed projectile and ballistic experts immersing themselves in trajectory shaping as a vital tool for energy management. Phillips (2008) vividly illustrated how composite guidance and altitude exploitation could conserve kinetic energy, extending range and elevating terminal speeds beyond the capabilities of traditional PNG profiles [4]. This reframing of "optimality" invited a broader perspective—not merely the quest for minimum miss or minimal control effort, but also the maximization of useful terminal energy, all while contending with the ever-changing realities of aerodynamics and atmospheric density.

As the late 2010s emerged, the spotlight shifted toward velocity maximization. Jeon (2020) explored the conditions under which terminal-speed-focused solutions closely resembled PNG-like behavior, particularly when $N \approx 3$ in simplified contexts. However, he also exposed the fragile nature of these insights as target movements and atmospheric variations shifted dramatically with altitude [5]. This revelation accentuated an enduring gap: classical PNG fails to account for energy awareness, and the assumptions supporting analytic optimality often crumble where operations become critical—grappling with maneuvering targets, steep gradients in drag and thrust, and the finite capabilities of actuators.

Entering the 2020s, a pragmatic synthesis emerged, where high-fidelity simulations were artfully employed offline, while lightweight policies executed seamlessly in real-time. He (2021) and Kim (2020) showcased the potential of learning-based guidance, utilizing offline optimization to train concise policies capable of delivering near-optimal accelerations with impressive onboard efficiency [6, 7]. Building on this innovative foundation, Zhang (2024) and Lee (2023) explored the realms of adaptive gain scheduling, hybrid PNG rules, and neural surrogates for managing complex logic, all designed for optimizing guidance around maneuvering targets and intricate non-convex envelopes [8, 9]. Most recently, Luo (2025) reported advancements in data-driven intercept-point prediction, implementing rapid gain adaptation that minimizes onboard computational demands while retaining accuracy in the face of target mobility and atmospheric variability [10]. Throughout these contributions, a clear and inspiring philosophy emerged: shift the heavy lifting of computation offline while keeping the flight software streamlined, agile, and certifiably reliable.

Yet, despite these remarkable advancements, two pivotal needs remain unfortunately unmet: (i) a midcourse guidance law that is explicitly aware of terminal speed and time of flight (TOF), which is interpretable and auditable—avoiding the pitfalls of black-box solutions, and (ii) a real-time mechanism capable of coupling TOF-consistent intercept predictions with energy-aware gain selection across the vast domain of altitude and downrange values. We propose an ambitious solution: constructing, offline, a discretized map of near-optimal PNG gains that try to maximize terminal speed under a constrained miss distance. During flight, this system intelligently retrieves the appropriate gain through interpolation based on a TOF-consistent intercept estimate. Compared to traditional PNG and time-optimal shaping [1–5], our approach maintains interpretability while weaving in an awareness of energy dynamics; when compared to contemporary learning-based methods [6–9, 11], it provides a comparable level of real-time efficiency enveloped in a transparent and verifiable lookup structure, smoothly allowing for integration with a compact regressor if the need arises.

2 Dynamic Model Description

We consider a planar missile–target engagement in which the missile is modeled as a point mass subject to kinematic guidance laws. The objective is to ensure negligible miss distance while maximizing the terminal impact velocity, thereby preserving sufficient kinetic energy at intercept. The engagement geometry is described by the missile and target positions

$$\boldsymbol{p}_m(t) = (x_m(t), z_m(t))^\top, \qquad \boldsymbol{p}_t(t) = (x_t(t), z_t(t))^\top, \tag{1}$$

from which the line-of-sight (LOS) angle

$$\lambda(t) = \arctan\left(\frac{z_t(t) - z_m(t)}{x_t(t) - x_m(t)}\right), \tag{2}$$

and its time derivative $\dot{\lambda}(t)$ are defined. The relative closing velocity is obtained as

$$V_c(t) = -\frac{d}{dt}\|\boldsymbol{p}_t(t) - \boldsymbol{p}_m(t)\|, \tag{3}$$

which together with $\dot{\lambda}(t)$ forms the basis of Proportional Navigation Guidance (PNG).

The commanded lateral acceleration under PNG is given by

$$a_c(t) = N\, V_c(t)\, \dot{\lambda}(t), \tag{4}$$

where N is the navigation constant. This parameter embodies the central trade-off of the guidance process. Small values of N induce shallow trajectories that conserve velocity but risk larger miss distances, while large values generate aggressive turns that suppress miss distance at the expense of higher energy loss and reduced terminal velocity. The challenge is therefore to select N so as to balance accuracy and energy retention.

We formalize this requirement as an optimization problem. Let $\mathcal{T} \subset \mathbb{R}^2$ denote the discretized set of admissible target states $\boldsymbol{p}_i = (x_i, z_i)$. For each state $\boldsymbol{p}_i$, we seek the navigation constant that maximizes the terminal impact velocity subject to a zero-miss constraint:

$$N^*(\boldsymbol{p}_i) = \underset{N \in [N_{\min}, N_{\max}]}{\arg\max} \; V_{\text{impact}}(N \mid \boldsymbol{p}_i), \tag{5}$$

$$d_{\text{miss}}(N \mid \boldsymbol{p}_i) \leq \varepsilon, \tag{6}$$

where V_{impact} is the missile's residual velocity at intercept, d_{miss} is the miss distance, and ε is a small tolerance approximating zero.

To solve this problem in practice, we discretize both the engagement space and the control parameter domain. The target domain is represented as

$$\mathcal{T} = \{\boldsymbol{p}_i = (x_i, z_i)^\top \mid i = 1, \ldots, M\}, \tag{7}$$

while the navigation constant is sampled from a structured lattice

$$\mathcal{N} = \{N_{\min} + k\,\Delta N \;\mid\; k = 0, 1, \ldots, K\}, \tag{8}$$

with ΔN denoting the discretization increment. This construction induces a full design space

$$\mathcal{D} = \mathcal{T} \times \mathcal{N}, \qquad |\mathcal{D}| = M \times (K + 1), \tag{9}$$

across which parametric simulations are executed.

For each $\boldsymbol{p}_i \in \mathcal{T}$, the simulations yield an optimal navigation constant $N^*(\boldsymbol{p}_i)$, the corresponding terminal velocity $V^*_{\text{impact}}(\boldsymbol{p}_i)$, and the associated time of flight $T^*(\boldsymbol{p}_i)$. Collecting these results over the entire domain produces a structured lookup table (LUT) that defines the mapping

$$\mathcal{F} : (x, z) \;\mapsto\; \big(N_{\text{opt}}(x, z),\; V^*_{\text{impact}}(x, z),\; T^*(x, z)\big), \tag{10}$$

During real-time operation, $\mathcal{F}$ provides a computationally efficient mechanism to retrieve the near-optimal gain, predicted terminal speed, and TOF without requiring online trajectory optimization. For query points that do not coincide with mesh nodes, bilinear or spline-based interpolation is employed.

The dynamic model described above therefore integrates the fundamental missile–target kinematics, the PNG law (4), and the optimization-based construction of an LUT for efficient real-time implementation. By embedding both accuracy and energy considerations directly into the guidance law design, this framework provides a balance between physical fidelity, computational tractability, and operational applicability.

3 Simulation Setup

To be able to grade performance of the proposed guidance algorithm targets with random velocity and orientation within predefined effective range. Algorithm can estimate interception point and choose scenarios where target stays within bounds. In nonlinear simulation domain engagement to these targets are performed and difference between estimated intercept point and simulated intercept point is compared.

3.1 Dynamical System Simulation Module

The nonlinear simulation environment is formulated as a deterministic dynamical system, where the complete aircraft dynamics are represented by a set of coupled, nonlinear, six degrees of freedom (6-DOF) equations of motion derived from the Newton–Euler formalism. Stochastic variations are not included, since optimization with respect to uncertain parameters is outside the scope of the present study. The system is implemented in a modular architecture that integrates atmosphere, aerodynamics, propulsion, mass properties, control system dynamics, and avionics subsystems into a unified simulation framework. The resulting model constitutes a highly accurate, nonlinear, time-domain representation.

The atmospheric model provides density, pressure, and temperature distributions as functions of altitude by employing MATLAB's `atmosisa` routine, which implements the International Standard Atmosphere (ISA). These quantities enter explicitly into both aerodynamic force/moment computation and propulsion performance equations.

Aerodynamic forces and moments are represented by a multidimensional lookup database generated via Computational Fluid Dynamics (CFD) simulations. The aerodynamic model provides nonlinear dependencies of force and moment coefficients on Mach number, altitude, angle of attack, and sideslip angle, along with dynamic derivatives for angular rate effects. This database ensures that the equations of motion capture stall phenomena, nonlinear cross-axis coupling, and unsteady aerodynamic effects essential for trajectory, stability and optimality analysis.

The propulsion system is modeled through a pressure–momentum formulation, in which thrust is computed from chamber pressure rise and nozzle pressure ratio relative to ambient conditions. This approach enables the accurate reproduction of thrust variation across the flight envelope, including transient phases such as throttle changes, accelerations, and climb segments.

The mass and inertia characteristics of the vehicle evolve during the mission due to fuel consumption and potential configuration changes. These effects are represented through time-varying mass properties accessed via a pre-computed database. The inertia tensor is updated accordingly, which directly influences rotational dynamics by changing coupling terms in the Newton–Euler equations.

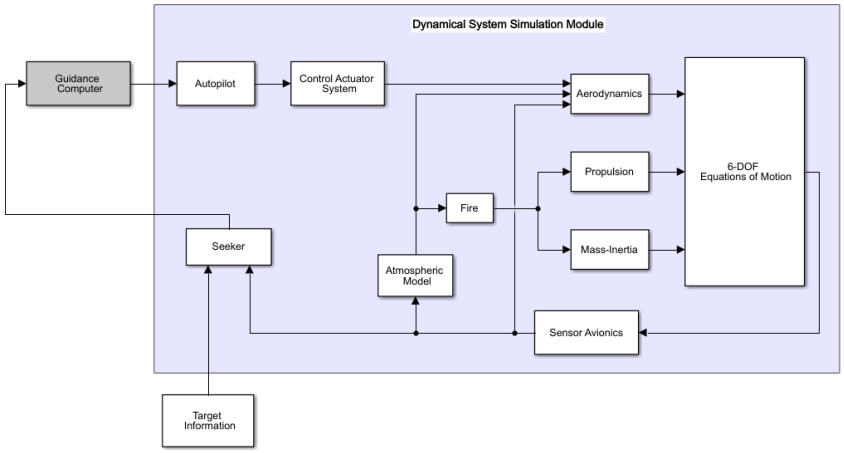

Fig. 1 Simulation Environment

The control system is modeled as a nonlinear dynamic subsystem, incorporating actuator dynamics, rate and deflection saturation, and transport delays. Actuators are treated as second-order systems with bandwidth constraints, ensuring that the closed-loop response reflects practical limitations rather than idealized control approach. The avionics subsystem is modeled with finite sampling, computational latencies, and sensor update rates, introducing additional delays into the control loop. These features are essential for realistic closed-loop stability margins and robustness analysis.

Numerical integration of the coupled nonlinear equations is carried out using fixed-step Runge–Kutta method. The modular structure of the simulation permits hierarchical data exchange between subsystems, emulating a real flight control system architecture.

By assembling atmosphere, aerodynamics, propulsion, mass and inertia variation, control dynamics, and avionics into the governing nonlinear differential–algebraic system, the complete 6-DOF model

is obtained. The simulation module thereby delivers a physically consistent platform for trajectory prediction, stability evaluation, and control system optimization under realistic operating conditions.

3.2 Experiments

Scenarios are generated by a random scenario generator that assigns targets with random velocity orientation, magnitude, and position. The algorithm verifies that each target remains within the missile's effective engagement envelope. Each scenario is then simulated in Simulink using a Runge–Kutta integrator with a 10^{-3} time step. If the miss distance falls below the selected threshold, the case is accepted as successful; the resulting time of flight and impact velocity are compared with the first-step estimates to assess the accuracy of the algorithm, since the offline solutions are already optimal within the discretization resolution.

The proposed guidance law is implemented in the model of a vertically launched missile with a two-stage thrust profile. The simulation switches to terminal guidance at a normalized distance of 0.25 from the target; the terminal phase uses a fixed PNG gain $N = 5$. During midcourse, the PNG gain is selected by the proposed algorithm. Three simulation cases are defined:

1) *Selection of Best Path*: For engagements directed at a single target point, the trajectory that minimizes the cost function is determined. This case provides insight into how the offline data of the algorithm are constructed.
2) *Single-Target Simulation*: The target is initialized at position $[1,\ 0.075]$ (normalized) with a speed of 0.0037 (normalized) and a flight-path angle of 135° measured from the positive downrange axis (i.e., climbing at 45° above the horizon toward the missile). The outputs of each guidance step are evaluated throughout the simulation.
3) *Multi-Target Performance*: A total of 48 equally spaced downrange–altitude grid points define the engagement domain. At each point, a target approaches with a normalized speed of 0.037 at constant altitude. The guidance law performance is evaluated across all 48 engagements.

4 Results

4.1 Selection of Best Path

For each intercept grid point, we perform a finite sweep of the PNG constant

$$\mathcal{N} \;:=\; \{\, N_k = 1.90 + 0.05\,k \;\mid\; k = 0, \ldots, 42 \,\} \quad (|\mathcal{N}| = 43), \tag{11}$$

and evaluate time of flight $T(N_k)$ and terminal speed $V_{\text{impact}}(N_k)$. Define the sweepwise extrema

$$V_{\max} = \max_{N\in\mathcal{N}} V_{\text{impact}}(N), \quad V_{\min} = \min_{N\in\mathcal{N}} V_{\text{impact}}(N), \qquad T_{\max} = \max_{N\in\mathcal{N}} T(N), \quad T_{\min} = \min_{N\in\mathcal{N}} T(N), \tag{12}$$

and the min–max–normalized performance indices (with a small $\eta > 0$ to avoid division by zero)

$$\widetilde{V}(N) \;:=\; \frac{V_{\max} - V_{\text{impact}}(N)}{V_{\max} - V_{\min} + \eta}, \qquad \widetilde{T}(N) \;:=\; \frac{T(N) - T_{\min}}{T_{\max} - T_{\min} + \eta}. \tag{13}$$

The user-weighted scalar objective is then

$$J(N) \;=\; w_V\,\widetilde{V}(N)^2 \;+\; w_T\,\widetilde{T}(N)^2, \qquad w_V, w_T \geq 0, \;\; w_V + w_T = 1, \tag{14}$$

and the selected gain is

$$N^\star \;\in\; \arg\min_{N\in\mathcal{N}} \; J(N). \tag{15}$$

Here, $V_{\max}, V_{\min}, T_{\max}, T_{\min}$ are taken over the same 43 simulations at the fixed grid point, so that $\widetilde{V}, \widetilde{T} \in [0, 1]$ are dimensionless and proportional.

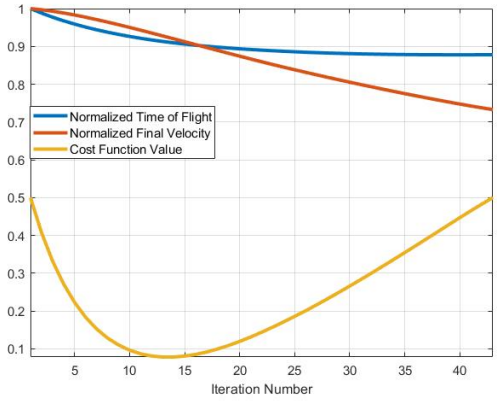

Fig. 2 Normalized terminal velocity $\widetilde{V}$, normalized TOF $\widetilde{T}$, and composite cost J versus sweep index k ($N_k = 1.90 + 0.05k$, $w_V = w_T = 0.5$).

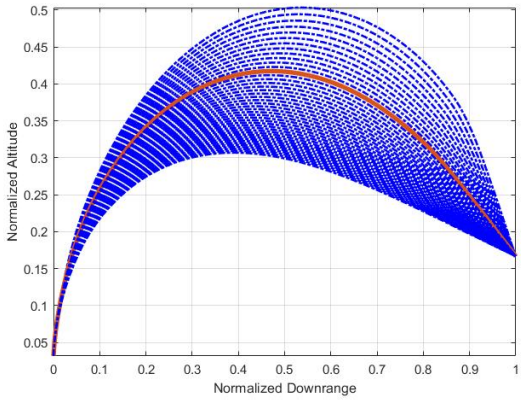

Fig. 3 Trajectories for different navigation constant

The normalized terminal-velocity and TOF traces for the 43 runs are summarized in Fig. 2. For smaller PNG gains (early samples of the sweep), both TOF and terminal speed are relatively large. Because the objective penalizes longer TOF while rewarding higher terminal speed, the composite cost decreases monotonically at first and then rises once further increases in N begin to erode the trade-off. The minimum occurs at the 13th sample of the grid (corresponding to $N = 2.50$), which is therefore adopted as the selected midcourse gain for this case.

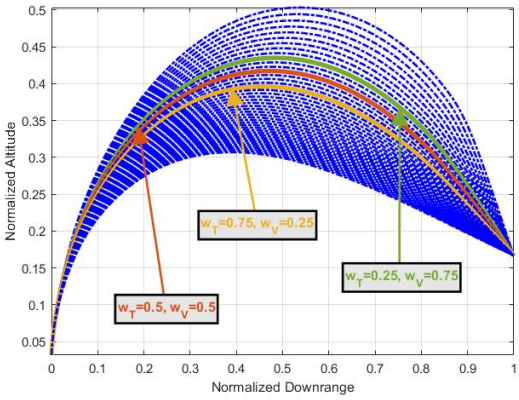

Fig. 4 Trajectories for different cost functions

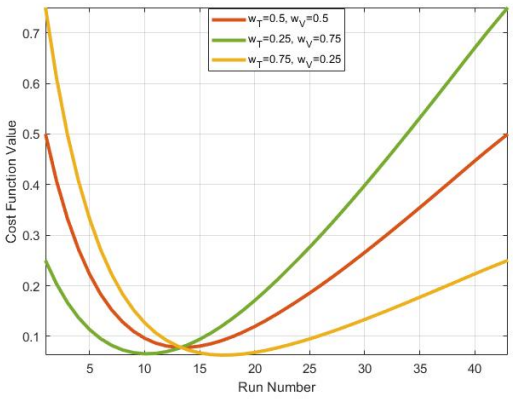

Fig. 5 Cost values for different cost functions

Figure 3 shows the trajectory of the 13th iteration compared to the runs with other PNG coefficients. It is evident that the 13th iteration neither has an excessively large TOF value nor sacrificing terminal speed—consistent with the minimum of the composite cost. By adding weights to the cost function (e.g., increasing w_T or w_V), an alternative path selection can also be made in line with the objective.

Figure 4 shows that when greater weight is assigned to final velocity, the missile, as expected, prefers to follow higher-altitude trajectories. In this case, midcourse guidance was carried out with $N = 2.35$, corresponding to the 10th run. On the other hand, when the weight of TOF is increased, the missile attempts to reach the target along a shorter path. Accordingly, during the 17th run, the midcourse phase directed the missile toward the target with higher PNG constant ($N = 2.70$).

4.2 Single-Target Simulation

Because the underlying dataset is proprietary, all quantities in Table 1 are reported in *normalized* units. Distances and speeds are scaled by the maximum downrange grid value, and times by the TOF at

Table 1 Comparison of Downrange Grids, TOF, Calculated Downrange, and Errors

Downrange Grids (d_i)	**TOF of Grids** (t_i)	**Calculated Downrange** (D_i)	**Error between Downrange Grids and Calculated Downrange** $(\Delta_i = d_i - D_i)$
0.65	0.6333	0.8201	-0.1701
0.70	0.6843	0.8056	-0.1056
0.75	0.7388	0.7902	-0.0402
0.80	**0.7901**	**0.7757**	**0.0243**
0.85	0.8459	0.7598	0.0902
0.90	0.8978	0.7450	0.1550
0.95	0.9501	0.7301	0.2198
1.00	1.0000	0.7160	0.2840

that point. Table 1 presents a representative subset of the database: the entries d_i (downrange grid) and t_i (corresponding TOF) are listed for uniformly spaced downrange samples with an increment of 0.025 at a fixed altitude $z = 0.075$.

Let the downrange–altitude axes be (x, z) and consider the fixed-altitude slice $z = 0.075$. The target state is initialized at

$$\boldsymbol{p}_t(0) = \begin{bmatrix} x_{t0} \\ z_{t0} \end{bmatrix} = \begin{bmatrix} 1 \\ 0.075 \end{bmatrix}, \qquad \boldsymbol{v}_t = \begin{bmatrix} v_t \cos\theta \\ v_t \sin\theta \end{bmatrix} = \begin{bmatrix} 0.0037 \cos 135^\circ \\ 0.0037 \sin 135^\circ \end{bmatrix}, \tag{16}$$

so the constant–velocity kinematic prediction is

$$\boldsymbol{p}_t(t) = \boldsymbol{p}_t(0) + \boldsymbol{v}_t\, t = \begin{bmatrix} x_{t0} + v_t \cos\theta\, t \\ z_{t0} + v_t \sin\theta\, t \end{bmatrix}, \tag{17}$$

For each grid TOF sample t_i in Table 1, we project the predicted target position onto the downrange axis to obtain the *candidate intercept downrange*:

$$D_i \;\coloneqq\; x_t(t_i) \;=\; x_{t0} + v_t \cos\theta\, t_i, \tag{18}$$

We then compare these continuous candidate values with the discrete downrange grid $\{d_i\}$ via the mismatch

$$\Delta_i \;\coloneqq\; d_i - D_i, \tag{19}$$

and select the index $i^\star = \arg\min_i |\Delta_i|$ as the closest grid intercept on the $z = 0.075$ slice. The pair $\left(d_{i^\star}, t_{i^\star}\right)$ seeds the subsequent LUT query and refinement.

To illustrate the method, we apply it to the scenario of Table 1. Scanning $\{|\Delta_i|\}$ yields the closest match at $d = 0.80$ with $|\Delta| = 0.0243$; the continuous target downrange there is $D = 0.7757$. Together with the target altitude at the corresponding time, this gives the initial intercept estimate $\boldsymbol{P}_0^1 = [\,0.7757,\ 0.2774\,]$.

The refinement proceeds as follows. At each iterate k, the LUT is queried at $\boldsymbol{P}_k$ to obtain the associated TOF t_k and optimal gain N_{opt}. The target is then forward-propagated to time t_k, yielding a projected intercept $\boldsymbol{P}_{\text{proj}}$, and the estimate is updated via $\boldsymbol{P}_{k+1} = \frac{1}{2}(\boldsymbol{P}_k + \boldsymbol{P}_{\text{proj}})$. For this example, the first query returns $t_0^1 = 0.6924$. After two updates the estimate changes only marginally, providing a stable seed for midcourse gain selection.

We monitor convergence by tracking the absolute change of a scalar progress metric between successive updates ($\Delta x_k = |x_k - x_{k-1}|$). If the change falls below a small threshold ε (i.e., the update is no longer materially improving the estimate), we stop; otherwise, we take another refinement step.

The numerical evidence in Tables 4, 5, and 6 in Appendix shows that the outputs at the 1st, 20,000th, and 40,000th simulation steps are essentially unchanged relative to the very first estimate. This indicates that the initial intercept/TOF seed is already near the fixed point. As the simulation progresses, a single refinement at a given step is typically sufficient to satisfy the stopping rule, making one-iteration updates adequate for both TOF and intercept estimation.

4.2.1 Final Velocity Data

Table 2 Altitude vs Downrange Values

Altitude \ Downrange	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95	1.00
0.175	0.0115	0.0114	0.0112	0.0113	0.0114	0.0114	0.0116	0.0117	0.0116
0.250	0.0119	0.0117	0.0116	0.0115	0.0115	0.0116	0.0116	0.0117	0.0116
0.325	0.0120	0.0117	0.0117	0.0116	0.0115	0.0115	0.0115	0.0113	0.0110

In Table 2, the final velocity values for scenarios simulated at different altitude and downrange values are shown. Based on this table, for the predicted interception point [0.8007 0.2743] obtained in the first step of the simulation, the estimated final velocity is 0.0114 (normalized). The estimated TOF from the first iteration is 0.7015 (normalized).

Table 3 Comparison of estimated and simulated values with relative error.

	First Step Estimated Value	Simulation Result	Relative Error
Final Velocity	0.0114	0.0116	1.7241%
TOF	0.7015	0.7007	0.1142%
Interception Point	[0.8007 0.2743]	[0.8010 0.2738]	[0.0375% 0.1826%]

In the first step of the simulation, the relative error values between the final velocity and TOF obtained from the offline database and those obtained from the executed simulation are as shown in Table 3.

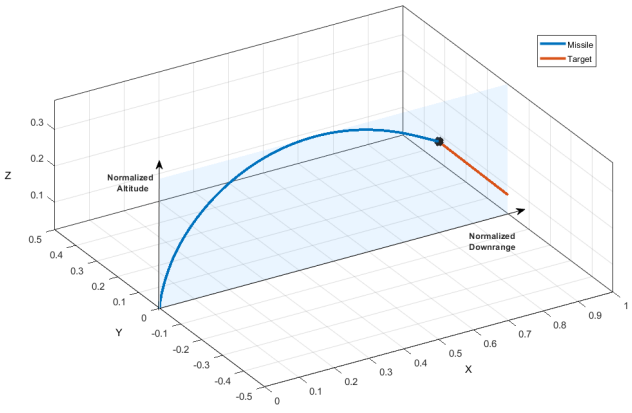

Fig. 6 Missile and target engagement trajectory (normalized (x, z) plane; parameters in text).

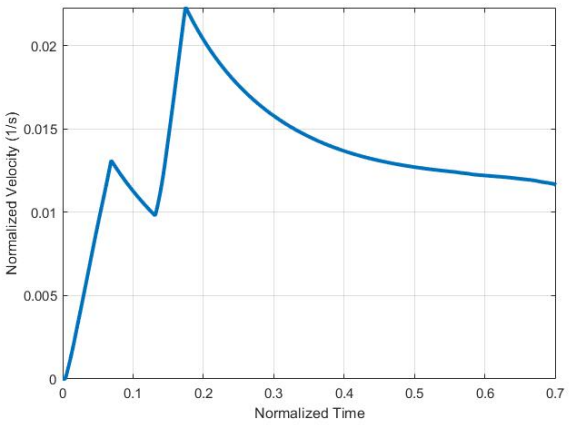

Fig. 7 Missile speed vs. time (normalized).

In Figures 6 and 7, the missile's trajectory and the velocity–time graph are shown. When the trajectory is examined, it can be seen that the missile ascends slightly above the interception altitude, and then enters the terminal phase to meet the target.

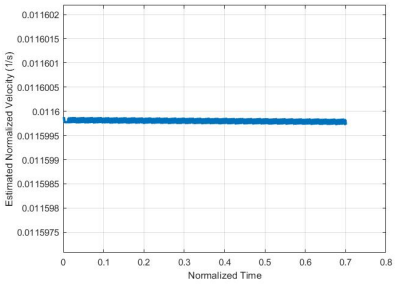

Fig. 8 Estimated impact velocity over time.

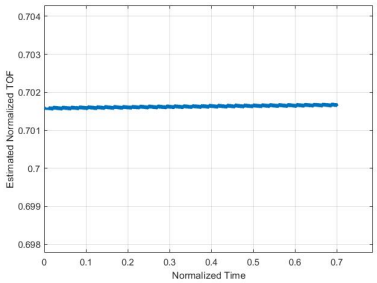

Fig. 9 Estimated TOF over time.

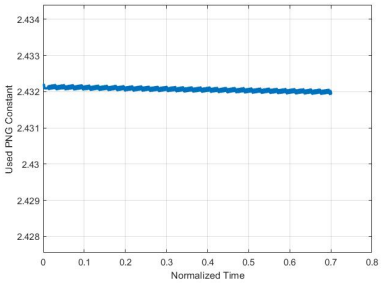

Fig. 10 Midcourse N over time.

In Figures 8, 9, and 10, the time-dependent output values of our guidance method during the simulation run are shown. It can be observed that the values at the beginning of the simulation remain *almost* unchanged until the end.

Figure 11 summarizes a 6×8 sweep over candidate interception points, comprising 48 closed-loop simulations. The grid is defined by six altitude levels (uniformly spaced across the admissible band) and eight downrange locations spanning the engagement corridor; all quantities are reported in the same normalized units used elsewhere. For each grid node, the midcourse gain is selected from the LUT using the TOF-consistent intercept estimate, and the terminal segment is flown with fixed PNG ($N = 5$).

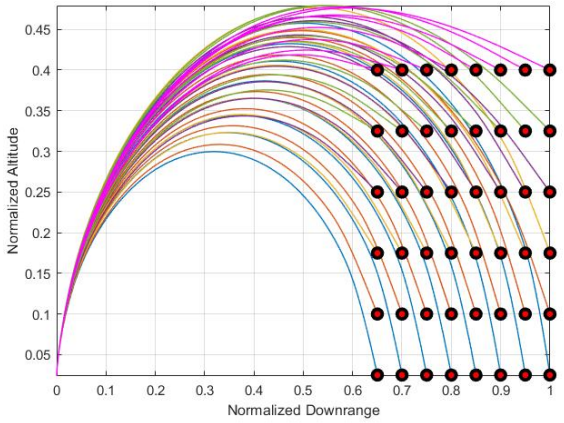

Fig. 11 Optimal trajectories for different intercept points

4.3 Multi-Target Performance

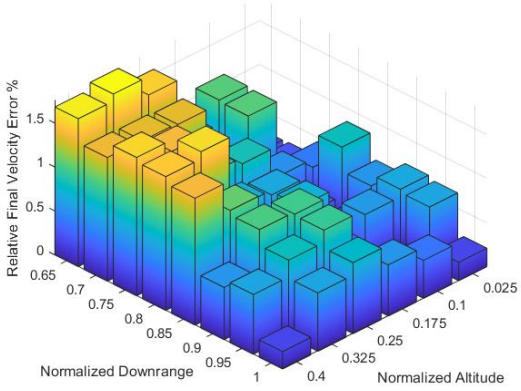

Fig. 12 Estimation error of impact velocity

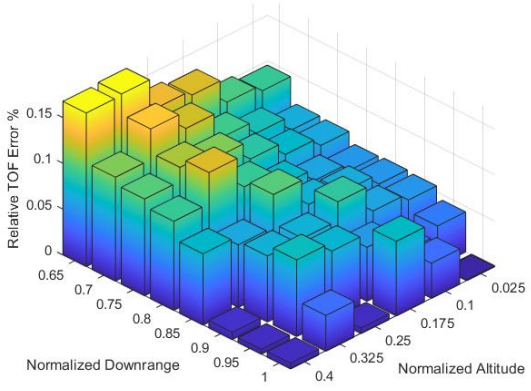

Fig. 13 Estimation error of time of flight

In Figures 12 and 13, the percentage relative errors of the final velocity and TOF values obtained from 48 runs, compared to the values predicted by the guidance method, are shown. The trends of the bar plots follow the same pattern. The highest error rates are observed at higher altitudes and shorter downrange distances. The guidance method described in this paper is applied for midcourse guidance, while the last 0.25 (normalized) of flight is conducted in the terminal phase with $N = 5$. For this reason, in cases where terminal guidance flight durations are short, the TOF and final velocity values are not strongly affected. The maximum percentage error in final velocity was found to be 1.762%, while for time of flight it was 0.1692%. These error levels appear to be within reasonable limits.

5 Conclusions and Future Work

This study presented a midcourse guidance strategy that couples a time–of–flight (TOF)–consistent intercept–point estimator with an offline, energy–aware lookup table (LUT) of near–optimal PNG gains. Evaluated across a 6×8 grid of interception geometries, the approach achieved sub-percent TOF error and $O(10^{-2})$ relative terminal-speed error (normalized units) while preserving the simplicity and transparency of PNG in real time. In effect, much of the "optimality gap" associated with fixed-gain PNG is closed by predicting a consistent intercept point early and retrieving a gain from an offline surface that already encodes altitude/downrange effects.

Empirically, the intercept estimate converges rapidly; after a short transient, a single refinement per guidance cycle is typically sufficient, keeping onboard computation modest. The LUT surface is smooth over the engagement domain, which enables stable interpolation without ad hoc filtering and yields coherent trajectories across neighboring grid points. When the terminal geometry remains close to the early prediction (e.g., constant-speed targets), the scheduled gain stays near its ex post optimum and the realized trajectory tracks the offline best case closely.

The main limitation appears when the predicted intercept shifts due to late target maneuvers or rapid geometry changes; gains applied prior to such shifts cannot be retrospectively corrected. A practical trade-off also exists between LUT resolution and memory footprint, and any model–plant mismatch in aerodynamics or propulsion can bias the precomputed surface if not recalibrated.

Future work will focus on strengthening robustness and broadening capability without sacrificing real-time efficiency. First, the LUT can be replaced or augmented by a compact supervised regressor trained on the offline dataset, reducing memory demands and smoothing extrapolation beyond the grid for embedded optimization. Second, a short-horizon maneuver predictor can be incorporated to hedge the scheduled gain against plausible near-term target actions, conditioning gain selection on a distribution of future intercepts rather than a single point estimate. Third, a unified midcourse–terminal co-design can remove phase discontinuities and jointly enforce impact-angle/time constraints alongside terminal-velocity objectives.

In summary, the proposed framework preserves the interpretability and low runtime cost of PNG while embedding energy awareness through offline computation. The outlined extensions target the remaining gap under maneuvering targets and model uncertainty, aiming toward a guidance architecture that is both high-performance and verifiable across a broad operating envelope.

Appendix

Table 4 Estimated TOF and Interception Point across Iterations (1st Step of Simulation)

Iteration	Estimated TOF	Estimated Interception Point
0	0.6924	[0.7976 0.2774]
1	0.6983	[0.7996 0.2754]
2	0.7004	[0.8003 0.2522]
3	0.7012	[0.8006 0.2745]
4	0.7014	[0.8007 0.2743]
5	**0.7015**	**[0.8007 0.2743]**
15	**0.7015**	**[0.8007 0.2743]**

Table 5 Estimated TOF and Interception Point for Two Iterations (2000th Step of Simulation)

Iteration	Estimated TOF	Estimated Interception Point
0	0.7015	[0.8007 0.2743]
1	**0.7015**	**[0.8007 0.2743]**

Table 6 Estimated TOF and Interception Point for Two Iterations (4000th Step of Simulation)

Iteration	Estimated TOF	Estimated Interception Point
0	0.7015	[0.8007 0.2743]
1	**0.7015**	**[0.8007 0.2743]**

Declaration of Use of Artificial Intelligence

The authors used Grammarly (including its Intelligent Assistance features) to improve the manuscript's language quality and clarity, and ChatGPT to assist in identifying recent and relevant sources during the literature review phase. These tools were used solely to increase research efficiency and support writing; all intellectual content, analysis, and conclusions are the authors' own. No AI system listed above contributed to the design, analysis, results, or interpretation of this work. Human oversight was applied to all AI-generated outputs, and full responsibility for accuracy and integrity of the manuscript rests with the authors.

References

[1] M. Guelman. Optimal guidance law in the plane. *Journal of Guidance, Control, and Dynamics*, 7(4):471–473, 1984. doi: 10.2514/3.19880.

[2] P. E. Kee, L. Dong, and C. J. Siong. Near optimal midcourse guidance law for flight vehicle. In *36th AIAA Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, 1998. AIAA Paper 1998-583. doi: 10.2514/6.1998-583.

[3] Neil F. Palumbo, Ross A. Blauwkamp, and Justin M. Lloyd. Modern homing missile guidance theory and techniques. *Johns Hopkins APL Technical Digest*, 29(1):42–59, 2010.

[4] Craig A. Phillips. Guidance algorithm for range maximization and time-of-flight control of a guided projectile. *Journal of Guidance, Control, and Dynamics*, 31:1447–1455, 2008. doi: 10.2514/1.31327.

[5] In-Soo Jeon, Mark Karpenko, and Jin-Ik Lee. Connections between proportional navigation and terminal velocity maximization guidance. *Journal of Guidance, Control, and Dynamics*, 43(2):383–388, Feb 2020. doi: 10.2514/1.G004672.

[6] Shaoming He, Hyo-Sang Shin, and Antonios Tsourdos. Computational missile guidance: A deep reinforcement learning approach. *Journal of Aerospace Information Systems*, 18(8):571–582, 2021. doi: 10.2514/1.I010970.

[7] Minjeong Kim, Daseon Hong, and Sungsu Park. Deep neural network-based guidance law using supervised learning. *Applied Sciences*, 10:7865, Nov 2020. doi: 10.3390/app10217865.

[8] Jiahui Zhang and Qiuqiu Wen. Closed-form guidance law for velocity maximization with impact angle constraint. *Journal of Systems Engineering and Electronics*, 35(5):1295–1303, 2024. doi: 10.23919/JSEE.2024.000078.

[9] Seungjae Lee, Jongho Shin, Hyeong-Geun Kim, Daesol Cho, and H. Jin Kim. Deep end-to-end imitation learning for missile guidance with infrared images. *International Journal of Control, Automation and Systems*, 21(10):3419–3429, 2023. doi: 10.1007/s12555-022-0683-6.

[10] Haowen Luo, Zichao Liu, Tianyu Jin, Chang-Hun Lee, and Shaoming He. Computational predictor–corrector homing guidance for constrained impact. *Journal of Guidance, Control, and Dynamics*, 48(6), 2025. doi: 10.2514/1.G008968.

[11] Hao Cui, K. Zhang, M. Tan, and J. Wang. A deep reinforcement learning-based cooperative guidance strategy under uncontrollable velocity conditions. *Aerospace*, 12(5):411, 2025. doi: 10.3390/aerospace12050411.