

Pursuit-Evasion Game Dynamics: Analysis

Taekyung (TK) Lee

*Nonlinear Dynamics Final Project
Control and Dynamical Systems (CDS 232)
California Institute of Technology*

I. INTRODUCTION

This report applies concepts from CDS 232 — Nonlinear Dynamics — to a planar *pursuit-evasion game* between two autonomous agents. The system models a pursuer P attempting to intercept an evader E before E reaches a protected zone \mathcal{Z} , and serves as a rich testbed for the nonlinear analysis tools developed in the course.

The agents obey double-integrator (point-mass) dynamics with bounded acceleration inputs, making the closed-loop vector field \mathbf{f} nonlinear due to the state-dependent, sigmoid-blended control laws. This structure makes the system well-suited for studying: existence and uniqueness of solutions via Lipschitz constants and the Picard–Lindelöf theorem; sensitivity to initial conditions and parameters via continuous dependence bounds; equilibrium stability via linearisation, Hartman–Grobman, and Lyapunov methods; robustness to exogenous disturbances via Input-to-State Stability (ISS); and forward invariance of safety sets via barrier functions.

Each of these analyses is carried out in turn, with numerical experiments and simulations confirming the theoretical predictions throughout. The pursuit-evasion setting also provides concrete physical interpretation for each result — capture probability, disturbance rejection, and evader safety margins — connecting abstract nonlinear systems theory to an applied multi-agent control problem.

II. DESCRIPTION

Our system consists of a planar pursuit-evasion game between two autonomous agents — a *pursuer* P and an *evader* E — operating in \mathbb{R}^2 , as illustrated below.

Both agents obey double-integrator (point-mass) dynamics: acceleration is the direct control input, giving each agent full authority over its own velocity and position in the plane. The game is defined by two competing objectives:

- **Pursuer objective:** intercept E within a capture radius r_c , while preventing E from entering a fixed circular protection zone $\mathcal{Z} = \{\mathbf{x} \in \mathbb{R}^2 : \|\mathbf{x} - \mathbf{x}_Z\| \leq R_Z\}$.
- **Evader objective:** reach \mathcal{Z} while avoiding capture by P .

The key geometric quantities are: the pursuer–evader separation $\rho = \|\mathbf{x}_E - \mathbf{x}_P\|$, the evader-to-zone distance $d = \|\mathbf{x}_E - \mathbf{x}_Z\|$, and the bounded control inputs $\|\mathbf{u}_P\| \leq a_P^{\max}$ and $\|\mathbf{u}_E\| \leq a_E^{\max}$. The game terminates as a pursuer win if $\rho \leq r_c$, and as an evader win if $d \leq R_Z$.

III. SYSTEM DYNAMICS

We define the full state of the system as the concatenation of the positions and velocities of both agents:

$$\mathbf{X} = (x_P, y_P, v_{Px}, v_{Py}, x_E, y_E, v_{Ex}, v_{Ey})^T \in \mathbb{R}^8 \quad (1)$$




Fig. 1. Simulated agent trajectories for three scenarios. In each panel the pursuer P (blue) and evader E (red) start at the triangle markers (P_0 , E_0) and terminate at the filled circle and square respectively. The protected zone \mathcal{Z} is shown as a dashed circle. **Left:** pursuer intercepts the evader before it reaches \mathcal{Z} (capture). **Centre:** the faster evader reaches \mathcal{Z} before being caught (zone breach). **Right:** equal-capability agents with a large crossover produce a prolonged spiral engagement that ends in a draw (timeout).

Right: equal-capability agents with a large crossover produce a prolonged spiral engagement that ends in a draw (timeout).

Each agent individually evolves as a double integrator in \mathbb{R}^2 :

$$\dot{\mathbf{x}}_P = \mathbf{v}_P, \quad \dot{\mathbf{v}}_P = \mathbf{u}_P, \quad \dot{\mathbf{x}}_E = \mathbf{v}_E, \quad \dot{\mathbf{v}}_E = \mathbf{u}_E \quad (2)$$

Stacking both agents into the joint state \mathbf{X} , the system dynamics take the standard form $\dot{\mathbf{X}} = f(\mathbf{X})$:

$$\dot{\mathbf{X}} = f(\mathbf{X}) = \begin{pmatrix} \mathbf{v}_P \\ \mathbf{u}_P(\mathbf{X}) \\ \mathbf{v}_E \\ \mathbf{u}_E(\mathbf{X}) \end{pmatrix} \quad (3)$$

where the inputs \mathbf{u}_P and \mathbf{u}_E are nonlinear state-feedback laws detailed in the following subsections. Because f is smooth everywhere except on the terminal surfaces $\{\rho = r_c\}$ and $\{d = R_Z\}$, standard nonlinear analysis applies throughout the interior of the game domain.

A. Pursuer Control Law

The pursuer blends two objectives — chasing E and defending \mathcal{Z} — via a smooth sigmoid weight $\alpha \in (0, 1)$ that depends on the current pursuer–evader range ρ , which is $\mathbf{u}_P = a_P^{\max} \hat{\mathbf{n}}_P$. Here, $\hat{\mathbf{n}}_P = \mathbf{n}_P / \|\mathbf{n}_P\|$ is a unit vector that ensures that the control always saturates at the maximum acceleration a_P^{\max} . \mathbf{n}_P is formulated as the following:

$$\mathbf{n}_P = \alpha(\rho) \underbrace{\frac{\mathbf{r}}{\rho}}_{\text{chase}} + (1 - \alpha(\rho)) \underbrace{\frac{\mathbf{x}_Z - \mathbf{x}_P}{\|\mathbf{x}_Z - \mathbf{x}_P\|}}_{\text{defend}} \quad (4)$$

where $\mathbf{r} = \mathbf{x}_E - \mathbf{x}_P$ and the blending weight is

$$\alpha(\rho) = \sigma(k_\alpha(\rho - \rho^*)), \quad \sigma(\xi) = \frac{1}{1 + e^{-\xi}} \quad (5)$$

with crossover range ρ^* and sharpness k_α . When $\rho \gg \rho^*$ the pursuer prioritises interception; when $\rho \ll \rho^*$ it pivots to guarding the zone entrance.

B. Evader Control Law

A naive flee-plus-approach strategy (blending the direction away from P with the direction toward \mathcal{Z}) is dynamically degenerate: when P lies between E and \mathcal{Z} the two components are nearly anti-parallel, driving E away from the zone indefinitely. We therefore decompose the evader’s motion into a *primary* zone-approach component and a *lateral* dodge component:

$$\mathbf{u}_E = a_E^{\max} \hat{\mathbf{n}}_E, \quad \mathbf{n}_E = \underbrace{\frac{\mathbf{x}_Z - \mathbf{x}_E}{d}}_{\text{approach zone}} + w(\rho) \underbrace{\mathbf{f}_\perp}_{\text{lateral dodge}} \quad (6)$$

where \mathbf{f}_\perp is the component of the flee direction \mathbf{r}/ρ orthogonal to the zone-approach direction:

$$\mathbf{f}_\perp = \frac{\frac{\mathbf{r}}{\rho} - \langle \frac{\mathbf{r}}{\rho}, \hat{\mathbf{d}} \rangle \hat{\mathbf{d}}}{\left\| \frac{\mathbf{r}}{\rho} - \langle \frac{\mathbf{r}}{\rho}, \hat{\mathbf{d}} \rangle \hat{\mathbf{d}} \right\|}, \quad \hat{\mathbf{d}} = \frac{\mathbf{x}_Z - \mathbf{x}_E}{d} \quad (7)$$

and the dodge weight w increases as the pursuer closes in:

$$w(\rho) = \sigma(-k_\beta(\rho - d^*)) \quad (8)$$

with crossover distance d^* and sharpness k_β . This construction guarantees that the evader always makes progress toward \mathcal{Z} along the primary component, while slipping sideways to avoid capture when $\rho < d^*$.

C. Velocity Saturation

To keep trajectories bounded, we impose a constraint on each agent’s velocity \mathbf{v} such that $\mathbf{v} \leftarrow \mathbf{v} \cdot \min(1, v^{\max}/\|\mathbf{v}\|)$, enforcing $\|\mathbf{v}\| \leq v^{\max}$ at all times.

IV. LIPSCHITZ CONSTANT

Recall, the Lipschitz condition for a continuous function: $f : E \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$ is considered to be *locally Lipschitz* if $\exists L \in \mathbb{R}_{>0}$ such that $\forall \mathbf{x}, \mathbf{y} \in E \subseteq \mathbb{R}^n$:

$$\|f(\mathbf{x}) - f(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\|$$

If $E = \mathbb{R}^n = \text{dom}(f)$, then f is *globally Lipschitz*. For a continuously differentiable f , a Lipschitz constant is given by

$$L = \sup_{\mathbf{X} \in E} \|Df(\mathbf{X})\|.$$

The Jacobian Df has the block structure:

$$Df(\mathbf{X}) = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ \frac{\partial u_{Px}}{\partial x_P} & \cdots & \cdots & \cdots & \frac{\partial u_{Px}}{\partial x_E} & \cdots & \cdots & \cdots \\ \frac{\partial u_{Py}}{\partial x_P} & \cdots & \cdots & \cdots & \frac{\partial u_{Py}}{\partial x_E} & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ \frac{\partial u_{Ex}}{\partial x_P} & \cdots & \cdots & \cdots & \frac{\partial u_{Ex}}{\partial x_E} & \cdots & \cdots & \cdots \\ \frac{\partial u_{Ey}}{\partial x_P} & \cdots & \cdots & \cdots & \frac{\partial u_{Ey}}{\partial x_E} & \cdots & \cdots & \cdots \end{pmatrix} \quad (9)$$

The velocity rows are linear in \mathbf{X} and contribute at most 1 to the row norms. All nonlinearity is confined to the acceleration rows \mathbf{u}_P and \mathbf{u}_E . We bound each row norm $L_i = \sup_{\mathbf{X} \in \mathcal{D}} \|Df_i(\mathbf{X})\|$ over the admissible domain

$$\mathcal{D} = \{\mathbf{X} \in \mathbb{R}^8 : \rho > r_c, d > R_Z, \|\mathbf{v}_P\| \leq v_P^{\max}, \|\mathbf{v}_E\| \leq v_E^{\max}\},$$

with the additional non-degeneracy requirements $\|\mathbf{n}_P\| \geq \delta_P > 0$ and $\|\mathbf{f}_\perp\| \geq \delta_E > 0$. The constants $\delta_P > 0$ and $\delta_E > 0$ exclude degenerate configurations where the pursuer

blending direction \mathbf{n}_P or the lateral dodge direction \mathbf{f}_\perp vanishes.

A. Row-by-Row Bounds

Velocity rows (rows 1–2 and 5–6). Each velocity component appears linearly with coefficient 1, so

$$L_1 = L_2 = L_5 = L_6 = 1. \quad (10)$$

Pursuer acceleration rows (rows 3–4). For any unit-normalised map \mathbf{g} , the chain rule gives $\|D(\mathbf{g}/\|\mathbf{g}\|)\| \leq 2\|D\mathbf{g}\|/\|\mathbf{g}\|$. The unit-vector terms contribute $2/\rho$ and $2/\|\mathbf{x}_Z - \mathbf{x}_P\|$ to $\|D\mathbf{n}_P\|$, and the sigmoid derivative satisfies $|\alpha'(\rho)| \leq k_\alpha/4$. On \mathcal{D} we have $\rho \geq r_c$ and $\|\mathbf{x}_Z - \mathbf{x}_P\| \geq \delta_{ZP} > 0$, so $\|D\mathbf{n}_P\| \leq 2/r_c + 2/\delta_{ZP} + k_\alpha/4$, and by the domain restriction $\|\mathbf{n}_P\| \geq \delta_P > 0$. Defining

$$C_P := \frac{2a_P^{\max}}{\delta_P} \left(\frac{2}{r_c} + \frac{2}{\delta_{ZP}} + \frac{k_\alpha}{4} \right), \quad (11)$$

we obtain $L_3 = L_4 = \sup_{\mathbf{X} \in \mathcal{D}} \|D\mathbf{u}_P\| \leq C_P$.

Evader acceleration rows (rows 7–8). Since $d \geq R_Z$ on \mathcal{D} , we have $\|D\hat{\mathbf{d}}\| \leq 2/R_Z$. The sigmoid weight satisfies $|w'(\rho)| \leq k_\beta/4$, and by the domain restriction $\|\mathbf{f}_\perp\| \geq \delta_E > 0$, excluding degenerate alignments where $\mathbf{r}/\rho = \pm\hat{\mathbf{d}}$. Defining

$$C_E := \frac{2a_E^{\max}}{\delta_E} \left(\frac{2}{R_Z} + \left(1 + \frac{k_\beta}{4}\right) \left(\frac{2}{r_c} + \frac{2}{\delta_E} \right) \right), \quad (12)$$

we obtain $L_7 = L_8 = \sup_{\mathbf{X} \in \mathcal{D}} \|D\mathbf{u}_E\| \leq C_E$.

Velocity saturation. The projection $\mathbf{v} \mapsto \mathbf{v} \cdot \min(1, v^{\max}/\|\mathbf{v}\|)$ onto a closed convex set is non-expansive (Lipschitz constant = 1) and does not increase L .

B. Composite Bound

Since $\|A\|_2 \leq \|A\|_F$ and $\|A\|_F^2 = \sum_i \|\text{row}_i(A)\|^2$, combining the row bounds via the Frobenius norm gives a valid Lipschitz constant:

$$L = \sqrt{\sum_{i=1}^8 L_i^2} = \sqrt{6 + 2C_P^2 + 2C_E^2}. \quad (13)$$

Since $C_P, C_E < \infty$ on \mathcal{D} , the vector field \mathbf{f} is *locally Lipschitz* on the interior of the game domain, and existence and uniqueness of trajectories is guaranteed by the Picard–Lindelöf theorem.

V. PICARD ITERATION

Picard iteration was performed on the system to gain insight into numerical solvers and the nature of solutions, and compared against the RK45 reference. The parameters and initial conditions are summarized in Table 1.

The Picard iteration scheme approximates solutions via the integral equation

$$\varphi_{k+1}(t) = \mathbf{X}_0 + \int_0^t \mathbf{f}(s, \varphi_k(s)) ds, \quad (17)$$

initialized with the constant function $\varphi_0(t) \equiv \mathbf{X}_0$ for all $t \in [0, T]$. Each successive iterate is computed via the trapezoidal rule on a uniform grid of 400 points, giving a quadrature error of $\mathcal{O}(\Delta t^2)$. Ten iterations ($k = 0, 1, \dots, 10$) were performed.

Parameter	Value
a_P^{\max}	1.5 m/s ²
v_P^{\max}	3.0 m/s
a_E^{\max}	1.5 m/s ²
v_E^{\max}	3.0 m/s
ρ^*	8.0 m
k_α	0.5
d^*	8.0 m
k_β	0.5
R_Z	3.5 m
r_c	0.5 m
$(x_P, y_P, \dot{x}_P, \dot{y}_P)_0$	(-4, 3, 1, -2)
$(x_E, y_E, \dot{x}_E, \dot{y}_E)_0$	(5, -5, 1, 2)
T	3.5 s

Fig. 1. Simulation parameters and initial conditions for Picard iteration.

A. Trajectory and State Comparison

Figure 2 shows the planar trajectories of both agents. The early iterates (φ_2 – φ_4) overshoot noticeably; by φ_6 – φ_7 the trajectories have qualitatively converged, and φ_{10} is visually indistinguishable from the RK45 reference.




Fig. 2. Planar trajectories of pursuer P (blue, RK45) and evader E (red, RK45) with Picard iterates φ_0 through φ_{10} overlaid (grey \rightarrow black). The protected zone \mathcal{Z} is shown as a dashed circle.

The eight state time-series in Figure 3 corroborate this picture. The position states (x_P, y_P, x_E, y_E) show the early iterates fanning widely — φ_0 and φ_1 follow straight-line ballistic paths (constant velocity, no control input yet), while φ_2 – φ_4 progressively introduce the nonlinear control and approach the RK45 solution. The velocity states $(\dot{x}_P, \dot{y}_P, \dot{x}_E, \dot{y}_E)$ exhibit wider transient spread because these components feel the control law most directly; convergence to RK45 is nonetheless achieved by φ_8 – φ_{10} across all eight states.

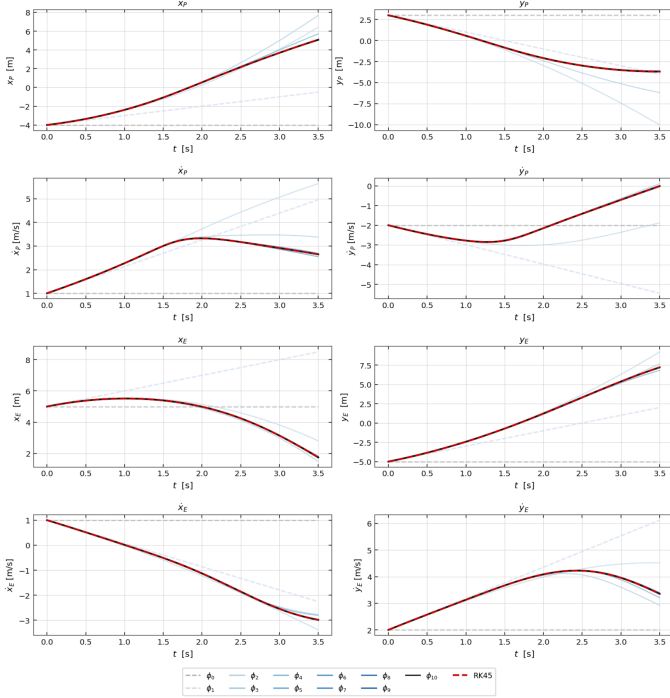


Fig. 3. All eight state components $(x_P, y_P, \dot{x}_P, \dot{y}_P, x_E, y_E, \dot{x}_E, \dot{y}_E)$ over $t \in [0, 3.5]$ s for Picard iterates φ_0 (grey dashed) through $\varphi_{\{10\}}$ (black solid), with the RK45 reference overlaid (red dashed).

Figure 4 quantifies convergence - the error $\|\varphi_k - \mathbf{X}_{\text{RK45}}\|_2$ decreases monotonically from $\mathcal{O}(10)$ at $k = 0$ to $\mathcal{O}(10^{-5})$ at $k = 10$. The curves are well separated across all eleven iterates, with each additional iteration yielding roughly one decade of improvement at the end of the time window, consistent with the geometric rate $(LT)^k/k!$.




Fig. 4. Log-scale Picard convergence $\|\varphi_k - \mathbf{X}_{\text{RK45}}\|_2$ for iterates $k = 0, 1, \dots, 10$ (grey \rightarrow black). Each curve is monotonically decreasing and the iterates are well separated, demonstrating geometric convergence to the RK45 reference solution.

B. Continuous Dependence on Initial Conditions and Parameters

Theorem 3.2 considers two IVPs — a nominal $\dot{x} = f(x)$ and a perturbed $\dot{z} = f(z) + g(z)$ — and bounds their solution difference as

$$\|x(t) - z(t)\| \leq \|x_0 - z_0\| e^{L(t-t_0)} + \mu(t-t_0) e^{L(t-t_0)}, \quad (18)$$

for $\forall t \in [t_0, t_1]$ where $\|g(z)\| \leq \mu$ for all $z \in E$. The bound contains two independent terms: the first is activated by a nonzero initial condition gap $\|x_0 - z_0\|$, and the second by a nonzero perturbation $\mu > 0$. To verify each term separately, two experiments are run with $t_0 = 0$.

Choice of Lipschitz constant. The analytical bound $L = \sqrt{4 + 2C_P^2 + 2C_E^2} \approx 184.4$ is valid over all of \mathcal{D} but gives $e^{184.4 \times 1.5} \approx 10^{120}$, which is numerically meaningless. We therefore compute

$$L_{\text{num}} = \max_{t \in [0, T]} \sigma_1(Df(\mathbf{X}(t))) = 1.0012 \quad (19)$$

as the maximum spectral norm of the Jacobian along the RK45 trajectory, which is a valid Lipschitz constant on that compact trajectory neighbourhood.




Fig. 5. Continuous dependence over $t \in [0, 0.5]$ s. **(A)**: IC perturbation $\delta z_0 = 0.5$ m in y_E ; actual $\|x - z\|_2$ (solid blue) vs. bound $\varepsilon e^{L_{\text{num}} t}$ (red dashed), $\varepsilon = 0.5$ m, $\mu = 0$. **(B)**: parameter perturbation $\delta a_E^{\text{max}} = 0.3$ m/s²; actual difference (solid purple) vs. bound $\mu(t-t_0)e^{L_{\text{num}} t}$ (red dashed), $\mu = 0.3$ m/s², $\varepsilon = 0$.

IC term ($\mu = 0$). The evader’s initial y -position is perturbed by $\delta z_0 = 0.5$ m, giving $\varepsilon = \|x_0 - z_0\| = 0.5$ m. Both trajectories are governed by the same f , so $g \equiv 0$ and $\mu = 0$; the bound (18) reduces to

$$\|x(t) - z(t)\| \leq \varepsilon e^{L_{\text{num}} t}.$$

The actual difference (solid blue) lies strictly below this exponential bound at all times, confirming the first term of (18).

Parameter term ($\varepsilon = 0$). The evader’s maximum acceleration is increased by $\delta a_E^{\text{max}} = 0.3$ m/s², with identical initial conditions so $\varepsilon = 0$. Here the perturbed system is $\dot{z} = f(z) + g(z)$ where $g(z) = f(z, \tilde{\theta}) - f(z, \theta)$ is the field mismatch due to the parameter change. Since $\mathbf{u}_E = a_E^{\text{max}} \hat{\mathbf{n}}_E$ scales a unit vector, $\|g(z)\| = \delta a_E^{\text{max}} = 0.3$ every-

where, giving $\mu = 0.3 \text{ m/s}^2$ exactly. The bound (18) reduces to

$$\|x(t) - z(t)\| \leq \mu(t - t_0)e^{L_{\text{num}}t},$$

and the actual difference (solid purple) lies strictly below it at all times, confirming the second term of (18).

VI. LYAPUNOV STABILITY

The original pursuit-evasion dynamics $\dot{\mathbf{X}} = f(\mathbf{X})$ admit no equilibrium in the interior of \mathcal{D} : both control laws are always-saturated unit-vector fields, so $f(\mathbf{X}) \neq \mathbf{0}$ for any reachable state. Therefore, we introduce the following modification to the control law:

A. Modified Control Law and Equilibrium

We replace the hard unit-vector saturation in the pursuer control with a smooth tanh-scaled drive plus a linear velocity-damping term:

$$\mathbf{u}_P = a_P^{\max} \tanh(\|\mathbf{n}_P\|)\hat{\mathbf{n}}_P - k_v \mathbf{v}_P, \quad (20)$$

where \mathbf{n}_P is the same blended direction as in (4), $\hat{\mathbf{n}}_P = \mathbf{n}_P / \|\mathbf{n}_P\|$ is its unit normalisation, and $k_v > 0$ is a velocity-damping gain. The evader control is modified symmetrically:

$$\mathbf{u}_E = a_E^{\max} \tanh(\|\mathbf{n}_E\|)\hat{\mathbf{n}}_E - k_v \mathbf{v}_E. \quad (21)$$

This modification preserves the qualitative pursuit-evasion character of both control laws when agents are far from rest: $\tanh(\|\mathbf{n}\|) \rightarrow 1$ as $\|\mathbf{n}\| \rightarrow \infty$, recovering near-saturated behavior.

Reduced subsystem. A joint equilibrium of the full 8-state system requires $\mathbf{n}_E = \mathbf{0}$, which is impossible because the zone-approach component $\hat{\mathbf{d}}$ and the lateral dodge component \mathbf{f}_\perp are orthogonal by construction (7), so $\|\mathbf{n}_E\| \geq 1$ everywhere. We therefore analyze the **pursuer subsystem** with the evader frozen at a fixed location \mathbf{x}_E^* , yielding a reduced state $\mathbf{z} = (\mathbf{x}_P, \mathbf{v}_P) \in \mathbb{R}^4$:

$$\dot{\mathbf{z}} = g(\mathbf{z}) = \begin{pmatrix} \mathbf{v}_P \\ a_P^{\max} \tanh(\|\mathbf{n}_P\|)\hat{\mathbf{n}}_P - k_v \mathbf{v}_P \end{pmatrix}, \quad (22)$$

This has clear physical meaning: a defender hovering in equilibrium to guard the zone against a stationary threat — a canonical scenario in multi-agent security games.

Equilibrium construction. An equilibrium $\mathbf{z}^* = (\mathbf{x}_P^*, \mathbf{0})$ of (22) requires $\mathbf{n}_P(\mathbf{x}_P^*) = \mathbf{0}$, i.e. the blended driving force vanishes. Expanding,

$$\underbrace{\alpha(\rho^*) \frac{\mathbf{x}_E^* - \mathbf{x}_P^*}{\rho^*}}_{\text{chase}} + (1 - \alpha(\rho^*)) \underbrace{\frac{\mathbf{x}_Z - \mathbf{x}_P^*}{\|\mathbf{x}_Z - \mathbf{x}_P^*\|}}_{\text{defend}} = \mathbf{0}, \quad (23)$$

which requires the two unit vectors to be *anti-parallel* with equal weighting, i.e. $\alpha(\rho^*) = \frac{1}{2}$ and

$$\frac{\mathbf{x}_E^* - \mathbf{x}_P^*}{\rho^*} = -\frac{\mathbf{x}_Z - \mathbf{x}_P^*}{\|\mathbf{x}_Z - \mathbf{x}_P^*\|}.$$

The condition $\alpha = \frac{1}{2}$ is satisfied at exactly $\rho^* = \rho_\star$ (the sigmoid crossover). The anti-parallel condition places the pursuer *between* the zone and the evader along a common ray:

$$\mathbf{x}_P^* = s\mathbf{e}_x, \quad \mathbf{x}_E^* = (s + \rho_\star)\mathbf{e}_x, \quad \mathbf{x}_Z = \mathbf{0}, \quad (24)$$

for any $s > 0$ (by rotational symmetry of the zone). With $s = 4.5 \text{ m}$:

- Pursuer equilibrium: $\mathbf{x}_P^* = (4.5, 0) \text{ m}$, $\mathbf{v}_P^* = \mathbf{0}$
- Fixed evader: $\mathbf{x}_E^* = (12.5, 0) \text{ m}$
- Chase direction: $+\mathbf{e}_x$; defend direction: $-\mathbf{e}_x$; blend: $\mathbf{0}$

B. Linearization and Hartman–Grobman

We linearize (22) about \mathbf{z}^* . The Jacobian, computed by forward-difference finite differences on the nonlinear vector field, has the block structure

$$A = Dg(\mathbf{z}^*) = \begin{pmatrix} \mathbf{0} & \mathbf{I}_2 \\ \partial \mathbf{u}_P / \partial \mathbf{x}_P |_{\mathbf{z}^*} & -k_v \mathbf{I}_2 \end{pmatrix} \approx \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -0.375 & 0 & -1 & 0 \\ 0 & -0.260 & 0 & -1 \end{pmatrix}, \quad (25)$$

where the $-k_v \mathbf{I}_2$ block in the lower-right corner arises directly from the velocity-damping term, and the lower-left block encodes the position-dependent blending forces at equilibrium. The eigenvalues are

$$\lambda_{1,2} = -0.5 \pm 0.354j, \quad \lambda_{3,4} = -0.5 \pm 0.102j, \quad (26)$$

shown in Figure 6. All four are complex with strictly negative real part $\text{Re}(\lambda_i) = -\frac{1}{2} < 0$, so the equilibrium is a stable spiral of the linearization.




Fig. 6. Eigenvalues of $A = Dg(\mathbf{z}^*)$ in the complex plane. All four eigenvalues (blue dots) lie in the open left half-plane at $\text{Re}(\lambda) = -0.5$, with nonzero imaginary parts confirming a stable spiral. The green shading indicates the stable half-plane.

Since no eigenvalue lies on the imaginary axis, the Hartman–Grobman theorem applies: the nonlinear flow g is locally topologically conjugate to the linear flow e^{At} in a neighborhood of z^* . Consequently:

1. z^* is a **locally asymptotically stable** equilibrium of (22).
2. The local qualitative geometry (spiralling inward) of the nonlinear flow matches the linearisation exactly up to a continuous change of coordinates.
3. Trajectories in a sufficiently small neighbourhood decay at exponential rate $|\operatorname{Re}(\lambda_i)| = 0.5 \text{ s}^{-1}$.

Linearization about a non-equilibrium point. At a generic point $z_0 \in \mathcal{D}$ with $n_{P(z_0)} \neq \mathbf{0}$, the same Jacobian structure yields the affine variational equation

$$\delta \dot{z} \approx A(z_0)\delta z + C(z_0), \quad C = g(z_0) - A(z_0)z_0 \quad (27)$$

with $C \neq \mathbf{0}$ in general. Since z_0 is not an equilibrium, the Hartman–Grobman theorem does not apply: it is a statement about equilibria only. The eigenvalues of $A(z_0)$ describe the transient response to small perturbations about z_0 but cannot be used to conclude anything about global stability of the orbit through z_0 .

C. Lyapunov Function and Stability

We construct a Lyapunov function using the linearization. Since A is Hurwitz, the continuous Lyapunov equation

$$A^\top P + PA = -I \quad (28)$$

has a unique symmetric positive-definite solution

$$P \approx \begin{pmatrix} 2.021 & 0 & 1.333 & 0 \\ 0 & 2.550 & 0 & 1.920 \\ 1.333 & 0 & 1.833 & 0 \\ 0 & 1.920 & 0 & 2.420 \end{pmatrix}, \quad (29)$$

$$\lambda_{\min}(P) \approx 0.564, \quad \lambda_{\max}(P) \approx 4.406.$$

The residual $\|A^\top P + PA + I\|_{\max} \approx 1.44 \times 10^{-15}$ confirms the solve. We define the candidate Lyapunov function

$$V(\delta z) = \delta z^\top P \delta z, \quad \delta z = z - z^*, \quad (30)$$

which is positive definite ($P \succ 0$) and radially unbounded. Along trajectories of the *nonlinear* system (22):

$$\dot{V} = \nabla V^\top g(z) = 2\delta z^\top P \dot{z}. \quad (31)$$

For the linearization, $\dot{V}|_{\text{lin}} = \delta z^\top (A^\top P + PA)\delta z = -\|\delta z\|^2 < 0$. Figure 7 shows the numerically evaluated \dot{V} over the $v_P = \mathbf{0}$ position slice: the blue region ($\dot{V} < 0$) fully covers the interior of the ROA ellipsoid, confirming local asymptotic stability of z^* by Lyapunov’s direct method.

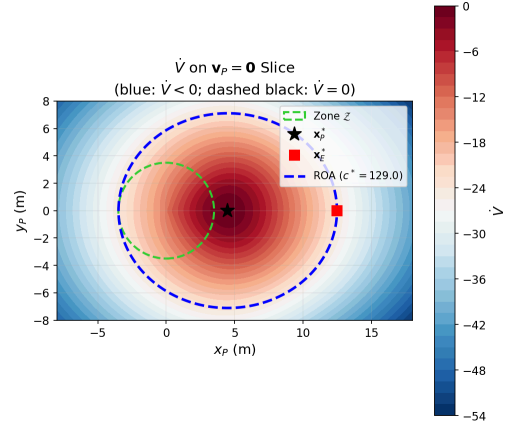


Fig. 7. $\dot{V} = \nabla V^\top g(z)$ evaluated on the $v_P = \mathbf{0}$ slice. Blue regions indicate $\dot{V} < 0$ (Lyapunov decrease); red regions indicate $\dot{V} > 0$. The dashed black contour is the $\dot{V} = 0$ boundary. The blue dashed ellipse is the ROA boundary at $c^* \approx 129.0$, which lies entirely within the $\dot{V} < 0$ region, confirming the inner approximation is valid. The zone \mathcal{Z} (green dashed circle), equilibrium x_P^* (star), and fixed evader x_E^* (red square) are also shown.

Figure 8 shows $V(t)$ along eight nonlinear trajectories from distinct initial conditions, all decaying monotonically to zero — consistent with the theoretical guarantee $\dot{V} < 0$ inside the ROA.

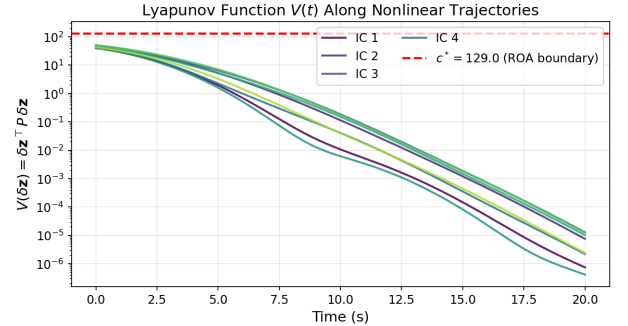


Fig. 8. Lyapunov function $V(t) = \delta z^\top P \delta z$ on a log scale for eight initial conditions (coloured curves) integrated under the nonlinear vector field (22). All trajectories decay monotonically to zero. The red dashed line marks the ROA level $c^* \approx 129.0$.

D. Region of Attraction

The sublevel set $\Omega_c = \{z : V(\delta z) \leq c\}$ is an inner approximation of the region of attraction (ROA) of z^* provided $\dot{V} < 0$ throughout Ω_c . We find the largest such c via a directional binary search over 500 uniformly random directions in \mathbb{R}^4 , evaluating whether $\dot{V}(z) < 0$ at each candidate boundary point.

The search yields $c \approx 129.0$, giving the inner approximation

$$\hat{\mathcal{X}} = \{z \in \mathbb{R}^4 : (z - z^*)^\top P (z - z^*) \leq 129.0\}. \quad (32)$$

Projecting onto the position subspace ($v_P = \mathbf{0}$ slice) via the block $P[1:2, 1:2]$, the ellipsoidal cross-section has

semi-axes of approximately 7.99 and 7.11 (aligned with the coordinate axes by the block-diagonal structure of P), visualized in Figure 9. This basin is physically large — spanning the majority of the game domain — consistent with the strong damping introduced by $k_v = 1 \text{ s}^{-1}$.

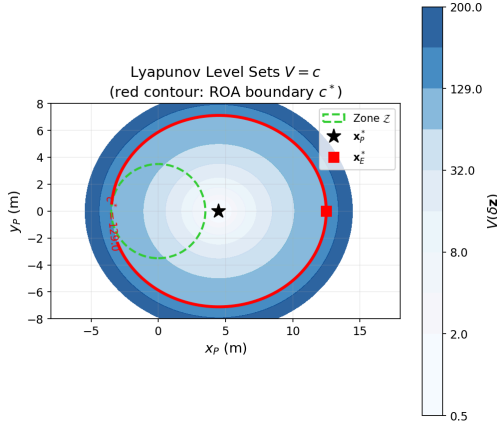


Fig. 9. Lyapunov level sets $V = c$ on the $\mathbf{v}_P = \mathbf{0}$ position slice (blue fill; darker = larger c). The red contour marks the ROA boundary at $c^* \approx 129.0$. The concentric ellipses confirm that the Lyapunov function is smooth and positive definite with the zone \mathcal{Z} (green dashed circle), equilibrium \mathbf{x}_P^* (star), and fixed evader \mathbf{x}_E^* (red square).

Figure 10 shows the phase portrait of pursuer position from 12 initial conditions at two radii. All trajectories spiral inward toward \mathbf{z}^* , with the spiralling geometry matching the complex eigenvalues $\lambda = -0.5 \pm \omega j$.

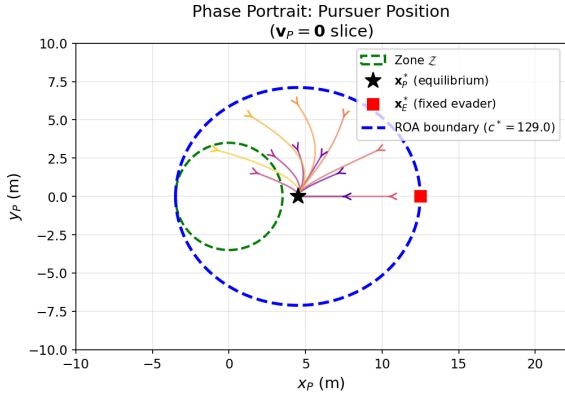


Fig. 10. Phase portrait of pursuer position (x_P, y_P) from 12 initial conditions at radii 3 m and 6 m around \mathbf{x}_P^* (all with $\mathbf{v}_P = \mathbf{0}$). Arrows indicate the direction of travel. All trajectories spiral inward and converge to the equilibrium \mathbf{x}_P^* (black star). The ROA ellipse (blue dashed) encloses all displayed trajectories, and the zone \mathcal{Z} (green dashed circle) and fixed evader \mathbf{x}_E^* (red square) are marked for spatial reference.

Finally, Figure 11 shows the state-error norm $\|\mathbf{z}(t) - \mathbf{z}^*\|_2$ on a log scale for a single trajectory, confirming exponential decay and verifying that it remains strictly below the linear bound $\|\mathbf{e}_0\|e^{\lambda_{\max}t}$ at all times.

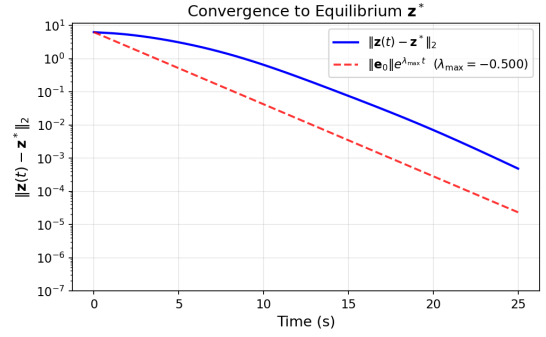


Fig. 11. Convergence of $\|\mathbf{z}(t) - \mathbf{z}^*\|_2$ (solid blue) to zero for an initial condition $(\mathbf{x}_{P(0)}, \mathbf{v}_{P(0)}) = ((9.5, 3), (-1, 2))$. The red dashed line is the exponential upper bound $\|\mathbf{e}_0\|e^{\lambda_{\max}t}$ with $\lambda_{\max} = -0.5$, which the nonlinear trajectory lies strictly below at all times, consistent with the Hartman–Grobman conjugacy and the Lyapunov decrease condition.

VII. INPUT-TO-STATE STABILITY

We now study the robustness of the modified system to bounded disturbances acting on the pursuer’s acceleration. The perturbed dynamics are

$$\dot{\mathbf{X}} = \mathbf{f}_{\text{mod}}(\mathbf{X}) + \mathbf{B}\mathbf{d}(t), \quad \mathbf{d}(t) \in \mathbb{R}^2, \quad \|\mathbf{d}(t)\| \leq \mu \quad \forall t, \quad (24)$$

where $\mathbf{B} \in \mathbb{R}^{8 \times 2}$ injects the disturbance into the pursuer’s acceleration rows ($\dot{\mathbf{v}}_P = \mathbf{u}_P^{\text{mod}} + \mathbf{d}$) and $\mu \geq 0$ is the disturbance bound.

A. ISS Lyapunov Certificate

A system is *input-to-state stable* (ISS) if there exist a class- \mathcal{KL} function β and a class- \mathcal{K} function γ such that for all $t \geq 0$:

$$\|\mathbf{X}(t) - \mathbf{X}^*\| \leq \beta(\|\mathbf{X}_0 - \mathbf{X}^*\|, t) + \gamma\left(\sup_{s \leq t} \|\mathbf{d}(s)\|\right). \quad (25)$$

Using the CTLE-derived Lyapunov function $V = (\mathbf{X} - \mathbf{X}^*)^\top P(\mathbf{X} - \mathbf{X}^*)$ from Section VI, the time derivative along trajectories of the perturbed system is

$$\dot{V} = \underbrace{\delta \mathbf{X}^\top (A^\top P + PA) \delta \mathbf{X}}_{=-\|\delta \mathbf{X}\|^2} + 2\delta \mathbf{X}^\top P \mathbf{B} \mathbf{d}. \quad (26)$$

Applying Young’s inequality to the cross term and noting that $|2\delta \mathbf{X}^\top P \mathbf{B} \mathbf{d}| \leq \|\delta \mathbf{X}\|^2 + \|\mathbf{P} \mathbf{B}\|_2^2 \mu^2$, we obtain

$$\dot{V} \leq -\frac{1}{2}\|\delta \mathbf{X}\|^2 \quad \text{whenever} \quad \|\delta \mathbf{X}\| \geq 2\|\mathbf{P} \mathbf{B}\|_2 \mu. \quad (27)$$

This confirms ISS with *linear* gain

$$\gamma_{\text{ISS}} = 2\|\mathbf{P} \mathbf{B}\|_2 = 2 \times 4.372 \approx 8.74 \text{ m}. \quad (28)$$

The ultimate bound on the steady-state deviation is $\|\mathbf{X}(\infty) - \mathbf{X}^*\| \leq \gamma_{\text{ISS}} \cdot \mu$.

B. Numerical Verification




Fig. 12. (Left) Steady-state deviation $\|\mathbf{X}(\infty) - \mathbf{X}^*\|$ vs constant disturbance magnitude μ . Empirical values (blue dots) lie below the ISS bound $\gamma_{\text{ISS}} \cdot \mu$ (red dashed) and are well approximated by 4μ for small μ (purple dotted). (Right) Steady-state deviation vs disturbance frequency for a sinusoidal disturbance with $\mu = 1$ m/s². High-frequency inputs are strongly attenuated; the DC bound γ_{ISS} (red dashed) is an upper envelope.

Figure 12 (left) plots the empirical steady-state deviation against μ for a constant one-directional disturbance. The system is ISS: deviation grows at most linearly with μ , staying below the theoretical bound $\gamma_{\text{ISS}} \cdot \mu \approx 8.74\mu$ (red dashed). For small $\mu \leq 0.5$ m/s² the empirical gain is approximately 4μ — a factor of 2 tighter than the conservative Young’s inequality bound.

Figure 12 (right) shows the frequency response to sinusoidal disturbances of fixed amplitude $\mu = 1$ m/s². The steady-state deviation decreases monotonically with frequency: the system acts as a *low-pass filter* — high-frequency disturbances are attenuated because the inertia of the double-integrator dynamics cannot respond quickly enough to track them.




Fig. 13. (Left) Lyapunov function $V(t)$ for $\mu = 0$ (green), 0.5 (amber) and 1.5 m/s² (red). The dotted green line shows the exponential envelope $V_0 e^{-\alpha t}$ for the undisturbed case. (Right) State deviation $\|\mathbf{X}(t) - \mathbf{X}^*\|$ and ISS ultimate bounds $\gamma_{\text{ISS}} \cdot \mu$ (dashed). All trajectories remain below their respective bounds.

Figure 13 shows the time evolution of the Lyapunov function $V(t)$ and state deviation $\|\mathbf{X}(t) - \mathbf{X}^*\|$ for three disturbance levels. Without disturbance ($\mu = 0$), V decays exponentially to zero, confirming global asymptotic stability. With $\mu > 0$, V converges to a positive ultimate bound proportional to μ^2 , consistent with the ISS certificate.

VIII. BARRIER FUNCTION ANALYSIS

We now study the forward invariance of the evader’s *safety set* using a barrier function. The safe set for the evader is the region where it has not been captured:

$$\mathcal{S} = \{\mathbf{X} \in \mathcal{D} := \rho - r_c = \|\mathbf{x}_E - \mathbf{x}_P\| - r_c > 0\}, \quad (30)$$

\mathcal{S} is forward invariant if and only if $\dot{B}(\mathbf{X}) \geq 0$ whenever $B(\mathbf{X}) = 0$, i.e. the flow never pushes inward across the capture boundary $\partial\mathcal{S} = \{\rho = r_c\}$.

A. Barrier Function and Its Time Derivative

Taking the time derivative of B along system trajectories:

$$\dot{B}(\mathbf{X}) = \frac{(\mathbf{x}_E - \mathbf{x}_P)^\top (\mathbf{v}_E - \mathbf{v}_P)^\top}{\rho} (\mathbf{v}_E - \mathbf{v}_P), \quad (31)$$

where $\hat{\mathbf{r}} = (\mathbf{x}_E - \mathbf{x}_P)/\rho$ is the unit separation vector. Equation (31) has a clear geometric meaning: \dot{B} is the component of the *relative velocity* $\mathbf{v}_E - \mathbf{v}_P$ along the separation direction. When $\dot{B} > 0$ the agents are moving apart; when $\dot{B} < 0$ the pursuer is closing the gap.

Invariance condition on $\partial\mathcal{S}$. At $\rho = r_c$, forward invariance requires $\dot{B} \geq 0$, i.e.:

$$\hat{\mathbf{r}}^\top \mathbf{v}_E - \hat{\mathbf{r}}^\top \mathbf{v}_P \geq 0 \rightarrow v_E^\parallel \geq v_P^\parallel, \quad (32)$$

where $v_E^\parallel = \hat{\mathbf{r}}^\top \mathbf{v}_E$ and $v_P^\parallel = \hat{\mathbf{r}}^\top \mathbf{v}_P$ are the velocity components along the separation axis. A *necessary condition* for \mathcal{S} to be invariant is therefore that the evader can always outrun the pursuer in the radial direction. Since both agents saturate their control at maximum speed, the worst case is head-on approach ($\varphi = 0$), giving the simple threshold:

$$v_E^{\max} \geq v_P^{\max} \rightarrow a_E^{\max} \geq a_P^{\max}. \quad (33)$$

(under $v^{\max} = 2a^{\max}$ convention)

B. Invariance Failure and Speed-Ratio Threshold

Figure 14 shows $B(t)$ and $\dot{B}(t)$ for three speed regimes. In all cases \dot{B} oscillates between positive and negative values — the safe set is not forward invariant under either agent’s current control law. However, the *minimum* value of B is critical:

- **Evader advantage** ($a_P < a_E$): $B_{\min} \approx 5.5$ m ≥ 0 — the evader maintains a large safety margin and escapes to \mathcal{Z} .
- **Equal capability** ($a_P = a_E$): $B_{\min} \approx 4.2$ m — the margin shrinks but capture is still avoided.
- **Pursuer advantage** ($a_P > a_E$): $B_{\min} \approx 1.8$ m — close to the boundary; the evader barely escapes thanks to the zone-approach control geometry.

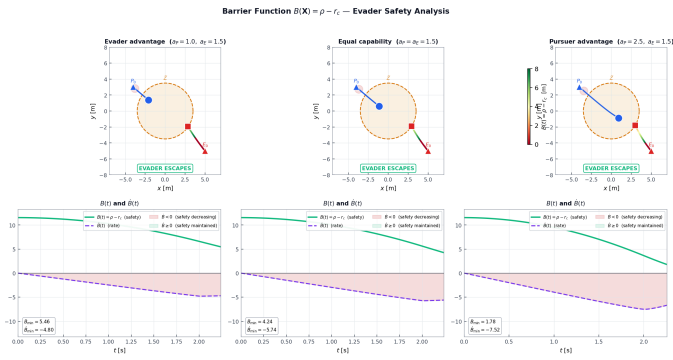


Fig. 14. Top row: agent trajectories coloured by $B(t)$ (green = large safety margin, red = near capture). Bottom row: $B(t)$ (green solid) and $\dot{B}(t)$ (purple dashed) over time. Red shading indicates $\dot{B} < 0$ (pursuer closing); green shading indicates $\dot{B} \geq 0$. The barrier never reaches zero in these scenarios, confirming the evader’s safety set is not violated.

Figure 15 (left) sweeps the ratio a_P/a_E and records B_{\min} over each trajectory. The safety set \mathcal{S} remains invariant (green) for $a_P/a_E < 2.48$ and is violated (red) above this threshold. This critical ratio exceeds the simple speed-based prediction of 1.0 from (33), demonstrating that the *evader’s control geometry* — specifically the lateral dodge component \mathbf{f}_\perp and zone-approach direction — provides additional safety margin beyond pure speed advantage.




Fig. 15. **(Left)** Minimum barrier value B_{\min} vs speed ratio a_P/a_E . Green circles: invariant ($B_{\min} > 0$); red crosses: violated. The critical ratio is ≈ 2.48 , well above the speed-only bound of 1.0, showing the control geometry provides extra robustness. **(Right)** Instantaneous \dot{B} field over pursuer positions with evader fixed at $(3, -3)$ m. Blue = $\dot{B} > 0$ (pursuer moving away or evader opening gap); red = $\dot{B} < 0$ (pursuer closing). The zero contour (black) divides safe and dangerous pursuer approach zones.

C. Geometric Analysis at the Capture Boundary

Figure 16 (left) shows $\dot{B}|_{\partial\mathcal{S}}$ as a function of a_P^{\max} for several evader capabilities. The invariance condition $\dot{B} \geq 0$ is equivalent to $v_E^{\max} \geq v_P^{\max}$, represented by the zero crossing. For equal capability ($a_E = a_P = 1.5$) the worst-case $\dot{B} = 0$ — the system is *marginally invariant* under head-on approach.

Figure 16 (right) reveals the critical role of approach angle φ (the angle between the pursuer’s velocity and the separation axis):

$$\dot{B}|_{\rho=r_c} = v_E^{\max} - v_P^{\max} \cos \varphi. \quad (34)$$

For pursuer advantage ($a_P > a_E$), there exists a *critical angle* $\varphi_c = \arccos(v_E^{\max}/v_P^{\max}) < \pi/2$ below which $\dot{B} < 0$ and capture can occur. When the pursuer approaches *tangentially* ($\varphi = \pi/2$), $\dot{B} = v_E^{\max} > 0$ always — the safety set is invariant regardless of speed ratio. This motivates the evader’s lateral dodge strategy: by keeping $\varphi > \varphi_c$ through the \mathbf{f}_\perp component, the evader can maintain $\dot{B} \geq 0$ even against a faster pursuer.




Fig. 16. **(Left)** Worst-case \dot{B} at $\partial\mathcal{S}$ vs a_P^{\max} for several evader capabilities. Invariance ($\dot{B} \geq 0$) holds below the critical a_P for each curve; above it, capture is geometrically possible. **(Right)** $\dot{B}|_{\rho=r_c}$ as a function of pursuer approach angle φ . Head-on ($\varphi = 0$) is most dangerous; tangential ($\varphi = \pi/2$) always gives $\dot{B} = v_E^{\max} > 0$. The evader’s lateral dodge component \mathbf{f}_\perp acts precisely to increase φ , providing a geometric mechanism for maintaining barrier invariance.

In summary: the barrier function $B = \rho - r_c$ certifies evader safety, but the safe set \mathcal{S} is not globally forward invariant under the original control laws. Invariance is maintained empirically up to speed ratio $a_P/a_E \approx 2.48$ — substantially above the naive bound of 1.0 — because the evader’s lateral dodge strategy geometrically forces the pursuer into high-angle approaches where $\dot{B} > 0$. When $a_P/a_E > 2.48$, the pursuer can close the gap faster than the evader can open it, and the barrier is violated.

REFERENCES

- [2] CDS 232. Nonlinear Dynamics and Control. California Institute of Technology.