Asymptotically Minimax Stochastic
Search Strategies in the Plane

by

Steven Lalley and Herbert Robbins
Purdue University and Columbia University

Technical Report #86-45

Department of Statistics
Purdue University

September 1986

Supported by NSF grant DMS 82-01723

◇

ABSTRACT

Stochastic search strategies are proposed for finding a possibly mobile target within a convex region of the plane. The strategies are asymptotically minimax as $\varepsilon \to 0$ with respect to the time required to get within ε of the target.

Asymptotically Minimax Stochastic
Search Strategies in the Plane

by

Steven Lalley and Herbert Robbins

1. Introduction

"Princess and Monster" (1) is a zero-sum two-person game with two players restricted to a bounded, connected, two-dimensional region Ω. The Monster (M) has maximum speed 1, the Princess (P) has maximum speed $v < 1$. Neither player obtains any information about the position of the other until the distance between the two is $\leq \varepsilon$; at which time M captures P and the game ends. The payoff to P is the time elapsed before capture.

This game is a crude model for a surface ship M attempting to locate a submarine. Here the parameter 2ε (the sweep width) is typically small relative to the dimensions of Ω.

The P and M game is too complex to admit simple minimax strategies. Even if the continuum Ω is replaced by a finite set of points, and even if P's strategy is known to M, M's optimal strategy can only be determined approximately, by a dynamic programming algorithm (2). Nevertheless, for convex Ω Gal (3,4) and Fitzgerald (5) have exhibited strategies for both players that are asymptotically minimax as $\varepsilon \to 0$, in the sense that the ratio of the expected payoff to the minimax value approaches 1 uniformly over the opponents' strategies. They have also shown that the minimax value $V(\varepsilon)$ satisfies

$$\lim_{\varepsilon \to 0} 2\varepsilon \, V(\varepsilon) = |\Omega|,$$

where $|\Omega|$ denotes the area of Ω.

Supported by NSF grant DMS 82-01723

P's strategy is easily described. Let $Q_1, Q_2, \dots$ be an i.i.d. sequence of random points uniformly distributed in Ω. P starts at A_1, stays there T time units, moves to Q_2 at full speed, stays there T time units, and so on. The parameter $T \to \infty$ as $\varepsilon \to 0$, but $\varepsilon T \to 0$ (e.g., $T = \varepsilon^{-1/2}$). It is not difficult to show that, no matter what strategy M uses, the expected time to capture is at least $|\Omega|/2\varepsilon$ (approximately) when ε is small.

M's strategy is more complicated. The region Ω is partitioned into long, narrow (width $\varepsilon^{1/2}$) rectangles. M searches in one of these rectangles for a long time T (e.g., $T = \varepsilon^{-1/3}$), then moves to another and searches in it for a time T, and so on (cf. Fitzgerald (4) or Gal (5) for details).

Despite its asymptotically minimax character, this strategy for M has a defect: when ε is small, M is confined to small subregions of Ω for very long periods of time. If the rules of the game were changed to allow P a small amount of partial information, e.g., if P were informed of the monster's position about once every $\varepsilon^{-1/3}$ time units, then she could elude it indefinitely. Thus, the Gal-Fitzgerald strategy for M is not robust to changes in the rules which might be relevant in naval operations.

In sections 2 and 3 below, we describe alternative strategies for M that are independent of ε, asymptotically minimax as $\varepsilon \to 0$, and are robust to changes in the rules allowing P occasional partial information.

The proofs of the main results will be published elsewhere. They depend on new methods for studying first passages to time-dependent boundaries by certain semi-Markov processes. The results of section 2 generalize results for search in a circle obtained by Lalley and Robbins (6).

2. Search in a Convex Region

Intuition suggests that a good strategy for M should produce random trajectories that are uniformly distributed in Ω, for if this were not the case, P could gain an advantage by hiding in the parts of Ω less frequently visited by M.

Let Ω be a bounded, convex region in $\mathbb{R}^2$ with smooth boundary $\partial\Omega$, and let ν be the normalized arc-length measure on $\partial\Omega$, $\int_{\partial\Omega} d\nu = 1$. Let $\Theta_1, \Theta_2, \ldots$ be i.i.d. random variables such that

$$P(\Theta_i \;\varepsilon\; d\theta) = \frac{1}{2} \sin\theta \; d\theta, \qquad 0 \leq \theta \leq \pi.$$

Define a sequence of (random) points $P_0, P_1, P_2, \ldots$ on $\partial\Omega$ as follows: Let P_0 have distribution ν. Having defined P_i, draw the chord in Ω from P_i that makes an angle Θ_{i+1} with the tangent to $\partial\Omega$ at P_i and define P_{i+1} to be the second point of intersection of the chord with $\partial\Omega$.

Proposition 1. The stochastic process $P_0, P_1, P_2, \ldots$ is a stationary, Harris-recurrent Markov chain on $\partial\Omega$ with stationary distribution ν.

See Revuz (7) for the definition of Harris-recurrence.

The trajectory of M is obtained by following the chords $P_0P_1, P_1P_2, \ldots$ in succession at unit speed. Let $X(t)$ denote the position of M at time $t \geq 0$.

Proposition 2. The stochastic process $X(t), t \geq 0$, is an ergodic semi-Markov process on Ω whose stationary distribution is the uniform distribution on Ω. In particular, if $f\colon\; \Omega \to \mathbb{R}$ is any continuous function, then

$$\lim_{t\to\infty} t^{-1} \int_0^t f(X(s)) \, ds = \int_\Omega f(\underset{\sim}{x}) \, d\underset{\sim}{x}/|\Omega| \quad \text{a.s.}$$

and

$$\lim_{t\to\infty} E(f(X(t))|X(0),\; \Theta_1) = \int_\Omega f(\underset{\sim}{x}) \, d\underset{\sim}{x}/|\Omega| \quad \text{a.s.}$$

Let Ω^0 denote the interior of Ω. For any $Q \varepsilon \Omega$ define

$$\tau_\varepsilon = \tau_\varepsilon(Q) = \inf\{t \geq 0: \ \text{dist}\ (X(t), Q) \leq \varepsilon\}.$$

Proposition 3. As $\varepsilon \to 0$,

$$2\varepsilon|\Omega|^{-1} E\ \tau_\varepsilon(Q) \longrightarrow 1$$

and

$$2\varepsilon|\Omega|^{-1}\tau_\varepsilon(Q) \xrightarrow{\mathcal{D}} \text{exponential with mean } 1,$$

uniformly for Q in any compact subset of Ω^0.

If P's strategy were to stay at a randomly chosen point of Ω, and if this were known to M, then M could do considerably better. By following an ε-dense path through Ω of approximate length $|\Omega|/2\varepsilon$, M could assure capture by time $|\Omega|/2\varepsilon$ (approximately) and the expected capture time would be (approximately) $|\Omega|/4\varepsilon$. Thus our plan is only 50% efficient for locating an immobile target.

For $\delta > 0$ let $\mathcal{F}_\delta$ denote the set of continuous, piecewise continuously differentiable functions $y(t), t \geq 0$, valued in Ω and such that $|y'(t)| \leq \nu < 1$ at all t where the derivative exists and dist $(y(t), \partial\Omega) \geq \delta$ for all t. For $y \varepsilon \mathcal{F}_\delta$ define

$$\tau_\varepsilon(y) = \inf\{t \geq 0: \ \text{dist}\ (X(t), y(t)) \leq \varepsilon\}.$$

Proposition 4. For any $\delta > 0$, as $\varepsilon \to 0$

$$\sup_{y \varepsilon \mathcal{F}_\delta} 2\varepsilon|\Omega|^{-1} E\tau_\varepsilon(y) \to 1.$$

This shows that our strategy of following the random trajectory $X(t)$ is almost asymptotically minimax. Since, according to the rules of the game, P is not required to stay at least δ away from $\partial\Omega$, M should actually follow the trajectory $X(t)$ for a convex region containing Ω in its interior, with some reasonable modification at $\partial\Omega$ (there is no point in searching outside Ω). It is clear that one may construct an asymptotically minimax family of strategies for M by using Proposition 4.

Our strategy does not have the "localization" property of the Gal-Fitzgerald strategy. Even if P is given the position and direction of M from time to time, she will not be able to predict its course for very long, in view of Proposition 2. Therefore, the strategy for M that we have described is not only fully efficient in the minimax sense, but also robust to partial information.

3. Search in a Parallelogram

A somewhat different search plan may be used when Ω is a parallelogram. Like that of the preceding section, this strategy does not suffer from the localization defect. For simplicity of exposition we shall describe the strategy for the square $\Omega = [0,1] \times [0,1]$.

Consider the torus $\hat{\Omega} = \mathbb{R}^2/2\mathbb{Z}^2$, and let π: $\mathbb{R}^2 \to \hat{\Omega}$ be the natural projection mapping. The torus $\hat{\Omega}$ may also be thought of as the square $[0,2] \times [0,2]$ with opposite sides identified. Consider the mapping ξ: $[0,2] \times [0,2] \longrightarrow [0,1] \times [0,1]$ defined by

$$\begin{aligned} \xi(x,y) &= (x,y) && \text{if } 0 \le x \le 1,\ 0 \le y \le 1; \\ &= (2-x,y) && \text{if } 1 \le x \le 2,\ 0 \le y \le 1; \\ &= (2-x,2-y) && \text{if } 1 \le x \le 2,\ 1 \le y \le 2; \\ &= (x,2-y) && \text{if } 0 \le x \le 1,\ 1 \le y \le 2. \end{aligned}$$

Since ξ maps corresponding points on opposite sides of $[0,2] \times [0,2]$ onto the same point of $[0,1] \times [0,1]$, ξ may be projected to a mapping ξ: $\hat{\Omega} \to \Omega$. The composition $\xi \circ \pi$: $\mathbb{R}^2 \to \Omega$ is a continuous mapping of the plane onto the unit square.

Fix $0 < \rho < 1$. Let $P_1, P_2, \ldots$ be i.i.d. random vectors in $\mathbb{R}^2$ with the uniform distribution on the circumference of the circle of radius ρ centered at $(0,0)$. Let P_0 be uniformly distributed on $[0,1] \times [0,1]$, and independent of $P_1, P_2, \ldots$. Define a random path $Y(t)$ in $\mathbb{R}^2$ as follows. Start at P_0, move at unit speed along the line segment from P_0 to $P_0 + P_1$, then move at unit speed along the line segment from $P_0 + P_1$ to $P_0 + P_1 + P_2$, and so on.

Our search plan calls for M to follow the projection onto Ω of the random path $Y(t)$, so that the position of M at time t is

$$X(t) = \xi(\pi(Y(t))).$$

Proposition 5. The stochastic process $X(t)$, $t \geq 0$, is an ergodic semi-Markov process on Ω whose stationary distribution is the uniform distribution on Ω. In particular, if $f\colon\ \Omega \to \mathbb{R}$ is continuous, then

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t f(X(s))\, ds = \int_\Omega f(\underset{\sim}{x})\, d\underset{\sim}{x} \quad \text{a.s.}$$

and

$$\lim_{t\to\infty} E(f(X(t))|P_0, P_1) = \int_\Omega f(\underset{\sim}{x})\, d\underset{\sim}{x} \quad \text{a.s.}$$

Let $\mathcal{F}_\delta$ be defined as in section 2. For $z \in \mathcal{F}_\delta$, let $\tau_\varepsilon(z) = \inf\{t \geq 0\colon\ \text{dist}\ (X(t), z(t)) \leq \varepsilon\}$.

Proposition 6. Let $0 < \delta < 1/2$. If $\rho < 2\delta$, then as $\varepsilon \to 0$,

$$\sup_{y \in \mathcal{F}_\delta} 2\varepsilon E\ \tau_\varepsilon(z) \to 1.$$

4. Concluding Remarks

(1) We conjecture that any ergodic semi-Markov process with uniform stationary distribution may be used to obtain asymptotically minimax strategies for M.

(2) The boundary effect implicit in Propositions 3, 4, and 6 deserves further study. How should M behave near $\partial\Omega$?

(3) It would be useful to have a good definition of a modified search game when the boundary $\partial\Omega$ is fuzzy.

(4) It would be useful to make a systematic study of modifications of the game in which one or both players is allowed partial information about the movement of the other.

References

(1) Isaacs, R. (1965) Differential Games (Wiley, New York).

(2) Brown, S. (1980) Opns. Res. 28, 1275–1289.

(3) Gal, S. (1979) SIAM J. Control Optim. 17, 99–122.

(4) Gal, S. (1980) Search Games (Academic, New York).

(5) Fitzgerald, C. (1979) SIAM J. Control Optim. 17, 700–712.

(6) Lalley, S. & Robbins, H. (1986) in Search Theory, ed. Chudnovsky (Springer, New York).

(7) Revuz, D. (1975) Markov Chains (North-Holland, Amsterdam).