A Geometric Strategy for Competitive Territory Expansion: The Largest Empty Circle Perspective

Lee Hao Rong Javier (as part of BT4015)

12/10/2025

1 Problem

Assume you are the *red player* in Figure 1, representing Shaw Theatres. Each red polygon shows your current geographic share within Singapore. The yellow polygons belong to your competitor, Golden Village. As an adversary to Golden Village, you wish to build a cinema (place a new dot) somewhere inside the Singapore boundary such that your *total Voronoi area* (the region closer to any red site than to any yellow site) becomes as large as possible. Where should this new site be placed to maximise your share?

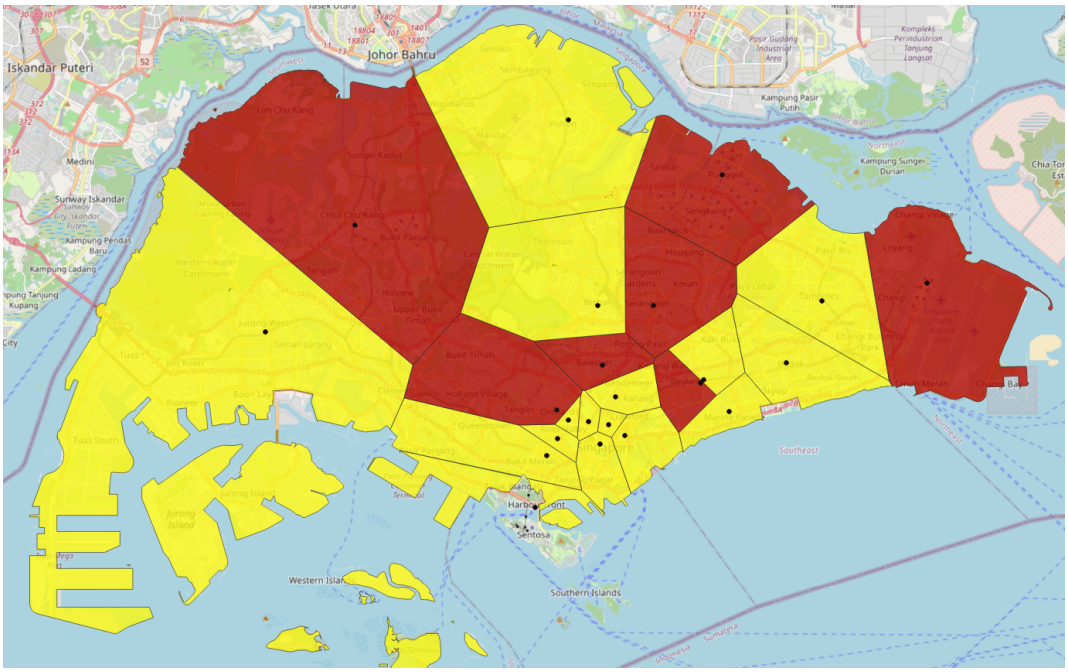


Figure 1: *Clipped Voronoi diagram of Singapore showing two competing site sets: $\mathcal{S}_R$ (red,* Shaw Theatres*) and $\mathcal{S}_Y$ (yellow,* Golden Village*). Each polygon $V_\Omega(s_i)$ is the locus within Ω closest to site s_i. The colour regions are $R_\Omega = \bigcup_{s_i \in \mathcal{S}_R} V_\Omega(s_i)$ and $Y_\Omega = \bigcup_{s_j \in \mathcal{S}_Y} V_\Omega(s_j)$.*

Formally, consider a finite set of sites $\mathcal{S} = \{s_1, s_2, \ldots, s_n\} \subset \mathbb{R}^2, n \in \mathbb{R}$ within a bounded polygonal domain $\Omega \subset \mathbb{R}^2$, such as the geographical boundary of Singapore. The *clipped Voronoi diagram* partitions Ω into regions

$$V_\Omega(s_i) = \{\, x \in \Omega : \|x - s_i\| \leq \|x - s_j\| \text{ for all } j \neq i \,\},$$

where each $V_\Omega(s_i)$ contains all points closer to s_i than to any other site, subject to the boundary constraint $x \in \Omega$. The dual structure is the Delaunay triangulation of $\mathcal{S}$ restricted to Ω, which provides geometric adjacency information between neighbouring sites.

Assume that the Voronoi cells are grouped into two disjoint colour classes:

$$\mathcal{S} = \mathcal{S}_R \cup \mathcal{S}_Y, \quad \mathcal{S}_R \cap \mathcal{S}_Y = \emptyset,$$

where $\mathcal{S}_R$ corresponds to *red* sites that are the geographic share of Shaw Theatres and $\mathcal{S}_Y$ to *yellow* sites that represent the geographic share of Golden Village. Let $R_\Omega = \bigcup_{s_i \in \mathcal{S}_R} V_\Omega(s_i)$ and $Y_\Omega = \bigcup_{s_j \in \mathcal{S}_Y} V_\Omega(s_j)$ denote the red and yellow subregions of the partition.

We want to introduce one additional site $p \in \Omega$, assigned the red label, such that the total red area after insertion:

$$A_R(p) = \text{area}(R_\Omega \cup V_\Omega(p)),$$

is maximised. Equivalently, we wish to maximise the *net area gain* to red:

$$p^\star = \arg\max_{p \in \Omega} \left[\text{area}(V_\Omega(p) \cap Y_\Omega) - \text{area}(V_\Omega(p) \cap R_\Omega)\right].$$

So the first term represents yellow territory captured by the new red site, while the second term penalises overlap with existing red regions.

We can draw parallels with the classical *Largest Empty Circle (LEC)* problem (Toussaint, 1983), where we find a point in Ω where its distance is maximised to both existing sites and the boundary:

$$p^* = \arg\max_{x \in \Omega} \min\left(\min_{s_i \in \mathcal{S}} \|x - s_i\|, \text{dist}(x, \partial\Omega)\right).$$

The Largest Empty Circle's centroid p^* is the location where a new site would have the largest possible Voronoi cell if colours were ignored. For our problem, the same geometric principles apply with the constraint that the point should preferentially occupy the yellow subregion Y_Ω, with the intent to enlarge the red area and avoid existing red dominance.

Therefore, we want to find a point that is in the yellow subregion $p^\star \in \Omega$ that balances:

1. the minimum distance of $p^\star$ from all existing sites and from boundary $\partial\Omega$ should be maximised.

2. $p^\star$ lies in or near yellow territory so $\Delta A_R(p)$ can be maximised.

1.1 Distance-to-constraints function

Define the *clearance radius* at any $x \in \Omega$ as

$$r(x) = \min\Big\{ \underbrace{\min_{1\le i\le n} \|x - s_i\|}_{\text{distance to nearest site}} \;,\; \underbrace{\operatorname{dist}(x, \partial\Omega)}_{\text{distance to boundary}} \Big\}.$$

Intuitively, $r(x)$ is the radius of the largest circle centred at x that lies fully within Ω and touches no existing site. The largest empty circle is thus

$$x^\star = \arg\max_{x\in\Omega} r(x), \qquad r^\star = \max_{x\in\Omega} r(x).$$

For the colour-constrained case, we can define a modified objective

$$r_R(x) = \begin{cases} r(x), & \text{if } x \in Y_\Omega, \\ \lambda\, r(x), & \text{if } x \in R_\Omega, \end{cases}$$

where $\lambda \in [0, 1)$ penalises points placed inside existing red regions. Then the optimal red insertion can be approximated by

$$p^\star = \arg\max_{x\in\Omega} r_R(x).$$

This is so that we can preserve the geometric structure of the LEC while encouraging placements that expand red coverage into yellow areas.

1.2 Lipschitz continuity and regularity

A function f is said to be L-Lipschitz if its rate of change is bounded by L everywhere:

$$|f(x) - f(y)| \le L\|x - y\|, \quad \forall x, y \in \Omega.$$

Every Lipschitz function is continuous, and hence at least one global maximum exists on any compact domain such as Ω. Both distance maps $x \mapsto \|x - s_i\|$ and $x \mapsto \operatorname{dist}(x, \partial\Omega)$ is 1-Lipschitz, and their minimum is also 1-Lipschitz. Therefore this means:

1. **There exists at least one global maximum and at least one global minimum.** Since Ω is a bounded and closed polygon and $r(x)$ is continuous, *extreme value theorem* guarantees that $r(x)$ has both a maximum and a minimum on Ω. Hence there exists at least one point where
$$x^\star = \arg\max_{x\in\Omega} r(x)$$
s.t. $r(x^\star) = r^\star = \max_{x\in\Omega} r(x)$. This means that the optimisation problem is well-defined and solvable.

2. **Controlled smoothness across Voronoi regions.** The Lipschitz bound
$$|r(x) - r(y)| \le \|x - y\|, \quad \forall x, y \in \Omega,$$

limits how fast the clearance radius can change as the centre moves. Thus, $r(x)$ cannot exhibit discontinuities or abrupt jumps when crossing Voronoi boundaries: moving the point by a small distance can change $r(x)$ by at most the same small amount (Preparata & Shamos, 1985). This bounded variation allows local gradient-based reasoning, since $r(x)$ varies smoothly almost everywhere in Ω.

3. **Differentiability almost everywhere.** By *Rademacher's theorem*, every Lipschitz function on $\mathbb{R}^n$ is differentiable almost everywhere (except on a set of measure zero). As a result, $r(x)$ possesses a well-defined gradient $\nabla r(x)$ at nearly all locations within Ω. This justifies my use of derivative-based arguments in subsequent results, where local stationarity conditions determine possible maximiser locations.

2 Where can the maximiser occur?

We maximise the red gain

$$J(x) = \operatorname{area}\big(V_\Omega(x) \cap Y_\Omega\big) - \operatorname{area}\big(V_\Omega(x) \cap R_\Omega\big), \qquad x \in \Omega,$$

Lemma 1 (No interior maximiser away from bisectors)**.** *Let $x \in \Omega$ be strictly inside a single Voronoi cell and away from $\partial\Omega$, and assume x is not coincident with any site. Then there exists a unit vector v and $\varepsilon > 0$ such that $J(x + tv) > J(x)$ for all $t \in (0, \varepsilon)$. Consequently, any local maximiser of J lies on a Voronoi edge or vertex, on $\partial\Omega$, or coincides with a site.*

Proof sketch with first–order geometry. Let s^* be the unique nearest site to x. The locus of points equidistant to x and s^* is the perpendicular bisector $\Pi(x, s^*)$. Move x by a small amount $t > 0$ along the unit normal v to $\Pi(x, s^*)$ pointing towards the yellow side if s^* is yellow, or away from the red side if s^* is red. Elementary geometry shows that the bisector translates by $\frac{t}{2}$ in the same normal direction:

$$\Pi(x + tv, s^*) = \Pi(x, s^*) + \tfrac{t}{2}v.$$

Let $L > 0$ be the length of the segment of $\Pi(x, s^*)$ that bounds $V_\Omega(x)$ inside Ω. Then the area of $V_\Omega(x)$ gained on the favourable side is $\frac{t}{2}\,L + o(t)$. All other boundaries of $V_\Omega(x)$ are not normal to v at x and their induced area change is $o(t)$. Hence the first variation of J at x in direction v is positive, contradicting local maximality. □

Lemma 2 (Finitely many stationary edge points)**.** *Let $\mathcal{B}$ be the finite arrangement of clipped Voronoi* edges and vertices *in Ω generated by the fixed sites $\mathcal{S}$. On any open Voronoi edge $e \subset \mathcal{B}$ the restriction $J\restriction_e$ is continuously differentiable and its stationary points are isolated. Together with finitely many Voronoi vertices, boundary intersections, and site locations, this yields a* finite *candidate set that contains every global maximiser of J.*

Why co–locating with a cinema can be optimal. Moving x inside a cell towards an existing site s pushes the bisector between x and s outwards and monotonically increases the newcomer's share of that cell until another constraint activates. In the limit $x \to s$ the two sites are coincident and consumers in that cell are indifferent between them, so no further first–order gain is possible by passing through s. Exact site locations are legitimate limit–candidates.

Take a yellow site $y \in \mathcal{S}_Y$ and choose a point $x_0 \in \operatorname{int} V_\Omega(y)$ at positive distance from every bisector and from $\partial\Omega$. Let n be the unit vector from x_0 to y and set $x_t = x_0 + tn$. For small $t > 0$ the only active boundary that moves at first order is the perpendicular bisector of $\overline{x_t y}$, which shifts away from y by $\frac{t}{2}$. If $L > 0$ denotes the length of the bisector segment bounding $V_\Omega(x_0)$ inside $V_\Omega(y)$, then

$$\operatorname{area}\big(V_\Omega(x_t) \cap Y_\Omega\big) = \operatorname{area}\big(V_\Omega(x_0) \cap Y_\Omega\big) + \tfrac{t}{2}L + o(t), \qquad \operatorname{area}\big(V_\Omega(x_t) \cap R_\Omega\big) = \operatorname{area}\big(V_\Omega(x_0) \cap R_\Omega\big) + o(t).$$

Therefore $J(x_t) - J(x_0) = \frac{t}{2}L + o(t) > 0$ for all sufficiently small $t > 0$. Hence x_0 cannot be a local maximiser. By symmetry, if x_0 lies strictly inside a red cell, moving x towards the nearest bisector reduces red overlap at first order and again increases J.

Every maximiser of J belongs to the finite set consisting of

1. Voronoi vertices in Ω (including boundary intersections) that separate at least two yellow cells;
2. stationary points of J along open Voronoi edges with at least two incident yellow cells (edge–wise first–order condition);
3. existing site locations (co–location limits).

Evaluating J on this finite set and choosing the maximiser yields the area–maximising insertion point.

Theorem 1 (Area–maximising coloured candidates)**.** *Let $p^\star \in \Omega$ maximise $J(x)$. Then $p^\star$ lies in the closure of Y_Ω and is one of:*

1. ***Shortest common-distance point on a Voronoi edge.*** *strictly inside Ω, with at least two of the incident cells yellow. Due to 1-Lipschitz, if there exists two points at the voronoi vertices that connect the edges that are extreme points, the straight line which extends the edge connecting these extreme points will include another stationary point and this point will occur at the shortest common-distance point on a Voronoi edge.*
2. ***Voronoi vertices*** *where a Voronoi edge intersects at least two Voronoi edge witihin boundary $\partial\Omega$, separating at least two or more yellow cells. Along each voronoi edge there exists a direction that strictly increases the amount of geographic share until a third constraint activates from at least one more site.*
3. ***Site locations*** *where you decide to compete with an initial cinema placement and place an additional point that is on the exact same point, such that users would be indifferent to go for either cinemas on the same point.*

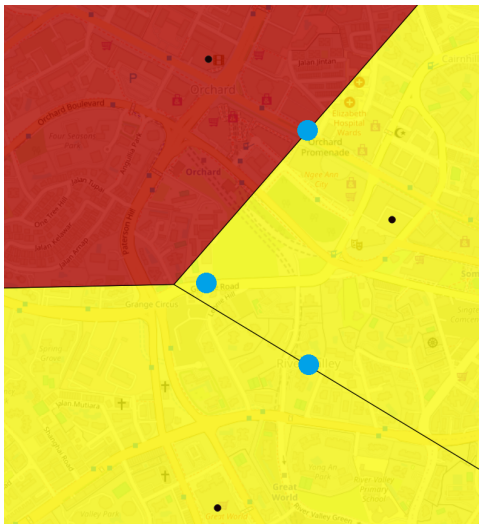


Figure 2: **Shortest common–distance point on a Voronoi edge.** The blue markers lie between the extension of two Voronoi edges strictly inside Ω.

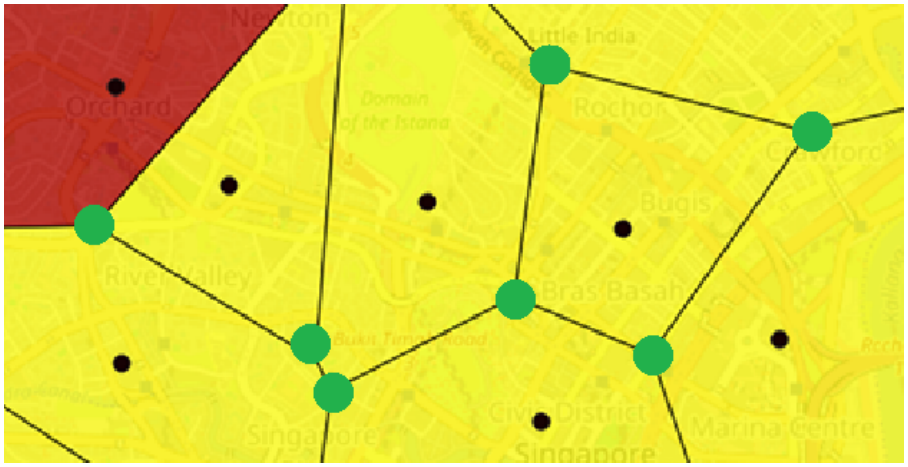


Figure 3: Voronoi vertices inside Ω: green markers show points where one Voronoi edge meets two other edges (edge–edge intersections). These are also the circumcentres of Delaunay triangles formed by the three nearest cinema sites (black dots).

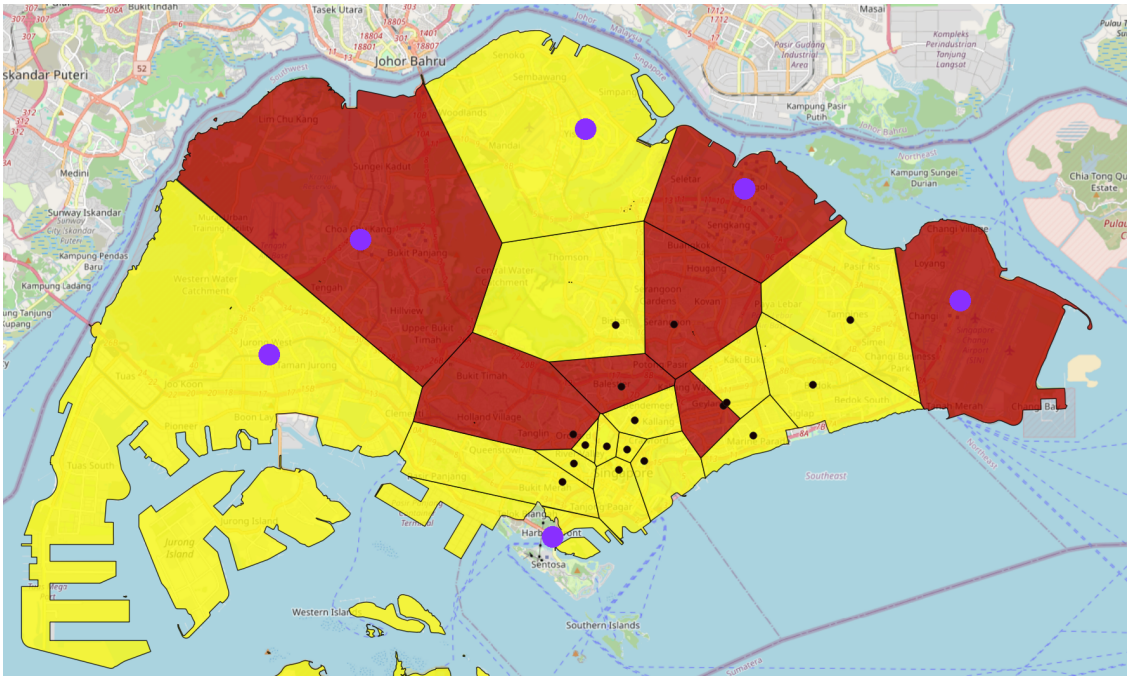


Figure 4: **Boundary Delaunay vertices** (purple): vertices of the Delaunay triangulation of the cinema sites (black points) that lie on the Delaunay convex hull. These vertices serve as boundary-dominant candidates when searching for red–area gains.

3 Algorithm

1. Compute the clipped Voronoi diagram of $\mathcal{S}$ in Ω and R_Ω, Y_Ω.
2. Enumerate coloured candidates (Theorem 1):
   (a) Shortest common-distance point on a Voronoi edge inside Ω between two yellow cells;
   (b) Voronoi vertices where at least two yellow cells are seperated;
   (c) Site locations with yellow sites.
3. For each candidate x, evaluate
$$J(x) = \text{area}\big(V_\Omega(x) \cap Y_\Omega\big) - \text{area}\big(V_\Omega(x) \cap R_\Omega\big).$$
4. Select $p^\star = \arg\max_x J(x)$ and insert it as the new red site.

4 Limitations toward Discrete Point Pattern Extensions

The current analysis relies crucially on the assumption that all sites form a finite, fixed configuration in a continuous Euclidean domain Ω, so that the clearance function
$$r(x) = \min\Big\{ \min_{1\le i\le n} \|x - s_i\|,\ \text{dist}(x, \partial\Omega) \Big\}$$

remains continuous and 1-Lipschitz. This ensures well-defined gradients, bounded variation, and the existence of at least one global maximiser and at least one global minimiser.

When one considers *discrete point pattern dynamics*—for example, stochastic insertion, deletion, or spatial resampling of sites as in spatial point processes or agent-based competition models—the Lipschitz property breaks down. Each infinitesimal perturbation in site position may create a topological change in the Voronoi diagram, causing discontinuous jumps in $r(x)$ and in the objective $J(x)$. In particular:

- The bisector network can reconfigure instantaneously when two sites become equidistant to a point x, producing non-differentiable shifts in the Voronoi boundaries.
- The function $r(x)$ becomes piecewise continuous with respect to site positions but no longer globally Lipschitz in either variable x or the site set $\mathcal{S}$.
- Gradient-based or geometric search methods (such as those based on the Largest Empty Circle) fail, since the continuity assumptions needed to identify stationary points do not hold.

References

[1] G. T. Toussaint, "The Largest Empty Circle Problem: Algorithms and Applications," *Pattern Recognition*, vol. 16, no. 4, pp. 315–325, 1983.

[2] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*. Springer, 1985.