Exploring Unknown Polygonal Environments with Discrete Visibility

Subir Kumar Ghosh

School of Technology & Computer Science
Tata Institute of Fundamental Research
Mumbai 400005, India

Overview

1. Robotic exploration problem
2. Visibility polygon
3. Exploring an unknown polygon: Continuous visibility
4. Exploring an unknown polygon: Discrete visibility
5. Exploring an unknown polygon: Bounded visibility

Robotic exploration problem

- Imagine that a robot is to explore the interior of a collapsed building, which has crumbled due to an earthquake, in order to search for human survivors.
- It is clearly impossible to have knowledge of the building's interior geometry prior to the exploration.
- Thus, the robot must be able to see, with its on-board vision sensors, all points in the building's interior while following its exploration path.
- In this way, no potential survivors will be missed by the exploring robot.

In our model for robotic exploration, we consider

- an unknown polygonal environment P with holes as the search space, and
- the robot as a moving point in P.

Visibility polygon

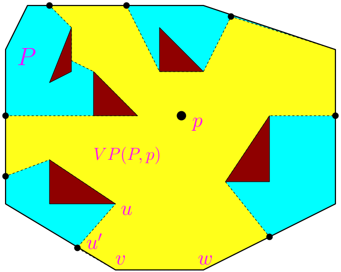

The visibility polygon of P from a point p (denoted as $VP(P,p)$) is the set of all points of P that are visible from p.

In other words, for every point $z \in P$, if the line segment joining z and p lies inside P, then z belongs to $VP(P,p)$.

In order to explore or see all points P, the robot computes visibility polygons from different locations inside P such that every point of P belongs to at least one such visibility polygon.

1. S. K. Ghosh, *Visibility Algorithms in the Plane*, Cambridge University Press, United Kingdom, 2007.

Continuous and discrete visibility

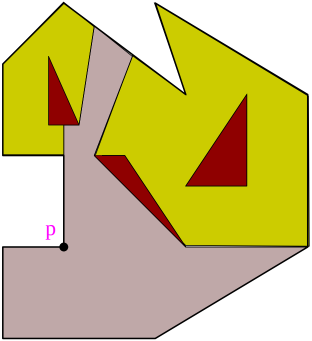

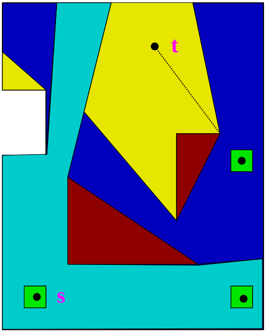

If the robot computes visibility polygons from each points on its path, we say that P is explored under continuous visibility.

If the robot computes visibility polygons from a selected set of points on its path, we say that P is explored under discrete visibility.

Efficiency of on-line algorithms

Suppose the polygon P is not known apriori and the point robot can compute the visibility polygon of P from its current position using visual sensors.

The robot wants to see all points of P with minimum cost. Cost can be the length or the number of links in the path that the robot has traveled starting from its initial position.

Efficiency of the on-line algorithm:

$$Competitive\ ratio = \frac{cost\ of\ \ the\ on-line\ algorithm}{cost\ of\ \ the\ off-line\ algorithm}$$

.

Algorithms for exploring an unknown environment with continuous visibility

1. A. Blum and P. Raghavan and B. Schieber, *Navigating in unfamiliar geometric terrain*, SIAM Journal on Computing, 26 (1997), 110-137.
2. K. Chan and T. W. Lam, *An on-line algorithm for navigating in an unknown environment*, IJCGA, 3 (1993), 227-244.
3. X. Deng and T. Kameda and C. Papadimitriou, *How to learn an unknown environment I: The rectilinear case*, Journal of ACM, 45 (1998), 215-245.
4. F. Hoffmann, C. Icking, R. Klein, K. Kriegel, *A competitive strategy for learning a polygon*, In Proceedings of the eight ACM-SIAM Symp. on Discrete Algo., pp. 166-174, 1997.
5. F. Hoffmann, C. Icking, R. Klein and K. Kriegel, *The polygon exploration problem*, SIAM Journal on Computing, 31 (2001), 577-600.
6. A. Lopez-Ortiz and S. Schuierer, *Searching and on-line recognition of star-shaped polygons*, Information and Computations, 185(2003), 66-88.

Searching for the kernel in an unknown polygon with continuous visibility

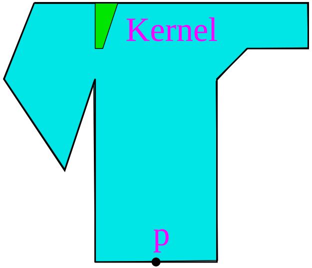

Starting from the initial position s, the problem is to design a competitive strategy to walk into the kernel of P.

Open Problem: The problem is open in link metric.

Algorithms for searching for the kernel with continuous visibility

1. C. Icking and R. Klein, *Searching for the Kernel of a Polygon—A Competitive Strategy*, SOCG, pages 258-266, 1995. Competitive ratio:5.331.
2. J.-H. Lee, C.-S. Shin, J.-H. Kim, S. Y. Shin and K.-Y. Chwa, *New competitive strategies for searching in unknown star-shaped polygons*, SOCG, pages 427-432, 1997. Competitive ratio: 3.828.
3. P. Anderson and A. Lopez-Ortiz, *A new lower bound for kernel searching*, CCCG, 2000. Competitive ratio: 1.515.

Searching for a target in a street with continuous visibility

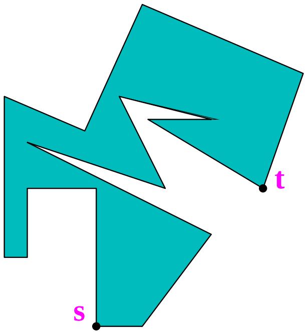

A *street* (also called *LR-polygon*) is a polygon for which two boundary chains from start to target are mutually weakly visible. So, the entire street is visible from any path from s to t.

The problem is to find a path from s to t such that the competitive ratio is the minimum.

1. S. K. Ghosh, *Visibility Algorithms in the Plane*, Cambridge University Press, 2007.

Algorithms for searching a street

1. R. Klein, *Walking an unknown street with bounded detour*, Computational Geometry: Theory and Applications, 1 (1992), 325-351. Competitive ratio: 5.72.
2. J. Kleinberg, *On line search in a simple polygon*, In Proceedings of the fifth ACM-SIAM Symposium on Discrete Algorithms, Pages 8-15, 1994. Competitive ratio: 2.83.
3. C. Icking, R. Klein, E. Langetepe and S. Schuierer, *An optimal competitive strategy for walking in streets*, SIAM Journal on Computing, 33(2004), 462-486. Competitive ratio: 1.41.
4. A. Lopez-Ortiz and S. Schuierer, *Lower bounds for streets and generalized streets*, International Journal of Computational Geometry and Applications, 11(2001), 401-421. Lower bounds: 1.41 and 9.06.
5. A. Datta and C. Icking, *Competitive searching in a generalized street*, CGTA, 13 (1999), 109-120. Competitive ratio: 9.06.
6. S. K. Ghosh and S. Saluja, *Optimal on-line algorithms for walking with minimum number of turns in unknown streets*, CGTA, 8 (1997), 241-266. Competitive ratio: 2.

Exploring unknown polygons: discrete visibility

In the remaining part of the lecture, we present exploration algorithms and their competitive ratios from the following papers.

1. S. K. Ghosh, J. W. Burdick, A. Bhattacharya and S. Sarkar, *On-line algorithms with discrete visibility: Exploring unknown polygonal environments*, Special issue on Computational Geometry approaches in Path Planning, IEEE Robotics and Automation Magazine, vol. 15, no. 2, pp. 67-76, 2008.
2. S. K. Ghosh and J. W. Burdick, *An on-line algorithm for exploring an unknown polygonal environment by a point robot*, Proceedings of the 9th Canadian Conference on Computational Geometry, pp. 100-105, 1997.
3. A. Bhattacharya, S. K. Ghosh and S. Sarkar, *Exploring an Unknown Polygonal Environment with Bounded Visibility*,Proceedings of the International Conference on Computational Science, Lecture Notes in Computer Science, No. 2073, pp. 640-648, Springer Verlag, 2001.

Motivation for discrete visibility

Many on-line computational geometry algorithms for exploring unknown polygons assume that the visibility region can be determined in a continuous fashion from each point on a path of a robot. Is this assumption reasonable?

1. Autonomous robots can only carry a limited amount of on-board computing capability.
2. At the current state of the art, computer vision algorithms that could compute visibility polygons are time consuming.
3. The computing limitations suggest that it may not be practically feasible to continuously compute the visibility polygon along the robot's trajectory.
4. For good visibility, the robot's camera will typically be mounted on a mast and such devices vibrate during the robot's movement.
5. Hence for good precision the camera must be stationary while computing visibility polygons.

It seems feasible to compute visibility polygons only at a discrete number of points.

Exploration cost

Is the cost associated with a robot's physical movement dominate all other associated costs?

The essential components that contribute to the total cost required for a robotic exploration can be analyzed as follows. Each move will have two associated costs as follows.

1. There is the time required to physically execute the move. If we crudely assume that the robot moves at a constant rate, r, during a move, the total time required for motion will be $r\ D$, where D is the total path length.
2. In an exploratory process where the robot has no apriori knowledge of the environment's geometry, each move must be planned immediately prior to the move so as to account for the most recently acquired geometric information. The robot will be stationary during this process, which we assume to take time t_M.
3. Since the robot is stationary during each sensing operation, we assume that it takes time t_S.

Let N_M and N_S be respectively the number of moves and the number of sensor operations required to complete the exploration of P. Hence, the total cost of an exploration is equated to the total time T required to explore P: $T(P) = t_M \ N_M + t_S \ N_S + r \ D$.

Now, $(t_M \ N_M + t_S \ N_S)$ can be viewed as the time required for computing and maintaining visibility polygons by computer vision algorithms, which is indeed a significant fraction of $T(P)$ because computer vision algorithms consume significant time on modest computers in a relatively cluttered environment.

Therefore, we assume that the overall cost of exploration is proportional to the cost for computing visibility polygons.

The criteria for minimizing the cost for robotic exploration is to reduce the number of visibility polygons that the on-line algorithms compute.

1. J. Borenstein and H. R. Everett and L. Feng, *Navigating mobile robots: sensors and techniques*, A. K. Peters Ltd., Wellesley, MA, 1995.
2. O. Faugeras, *Three-dimensional computer vision*, MIT Press, Cambridge, 1993.

An exploration algorithm

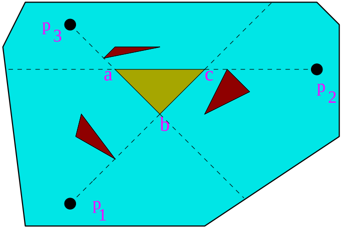

- We present an exploration algorithm that a point robot can use to explore an unknown polygonal environment P under discrete visibility.
- In order to explore P, the robot starts from a given position, and sees all points of the free space incrementally.
- It may appear that it is enough to see all vertices and edges of P in order to see the entire free-space. However, this is not the case.
- Three views from p_1, p_2 and p_3 are enough to see all vertices and edges of P but not the entire free-space of P.

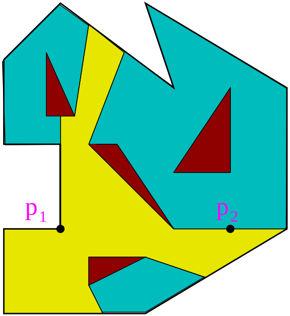

(i) Let S denote the set of viewing points that the algorithm has computed so far. (ii) The triangulation of P is denoted as $T(P)$. (iii) The visibility polygon of P from a point p_i is denoted as $VP(P, p_i)$.

Step 1: $i := 1$; $T(P) := \emptyset$; $S := \emptyset$; Let p_1 denote the starting position of the robot.
Step 2: Compute $VP(P, p_i)$; Construct the triangulation $T'(P)$ of $VP(P, p_i)$; $T(P) := T(P) \cup T'(P)$; $S = S \cup p_i$;
Step 3: *While* $VP(P, p_i) - T(P) = \emptyset$ *and* $i \neq 0$ *then* $i := i - 1$;

Step 4: *If* $i = 0$ *then goto* Step 7;
Step 5: *If* $VP(P, p_i) - T(P) \neq \emptyset$ *then* choose a point z on any constructed of $VP(P, p_i)$ lying outside $T(P)$;
Step 6: $i := i + 1$; $p_i := z$; *goto* Step 2;
Step 7: Output S and $T(P)$;
Step 8: Stop.

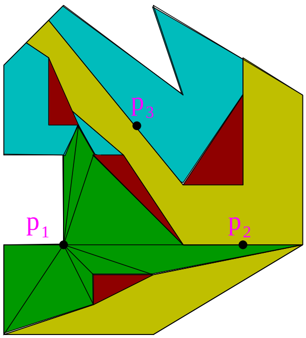

Competitive ratio

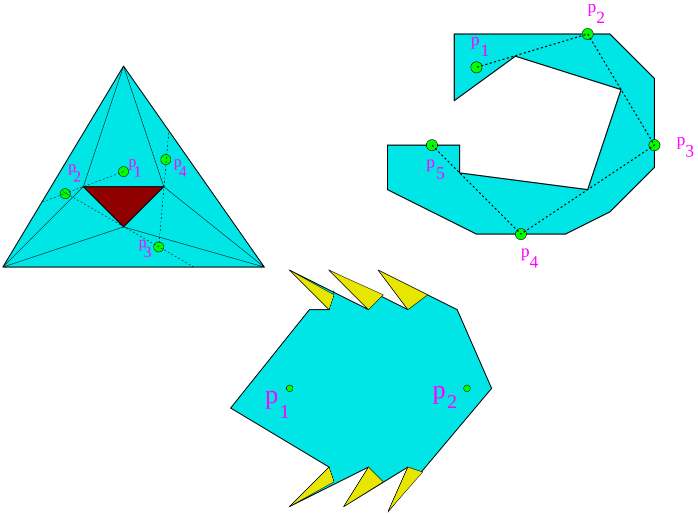

The algorithm needs $r + 1$ views. Competitive ratio is $(r + 1)/2$, where r denotes the number of reflex vertices of the polygon.

Open Problem: Can the bound be improved?

Convex robot exploration

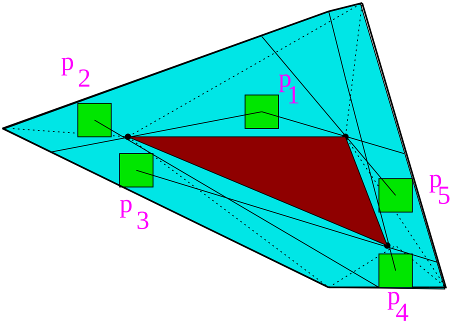

We wish to design an algorithm that a convex robot C can use to explore an unknown polygonal environment P (under translation) following the similar strategy of a point robot.

C needs more than $r + 1$ views for exploration.

Open problem

Can one derive an upper bound on the number of views for a convex robot exploration?

The art gallery problem

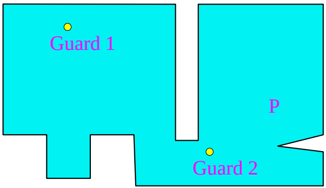

An art gallery can be viewed as a polygon P with or without holes with a total of n vertices and guards as points in P.

Victor Klee asked in 1976: How many guards are always sufficient to guard any polygon with n vertices?

The minimum vertex, point and edge guard problems for polygons with or without holes (including orthogonal polygons) are NP-hard.

1. J. O'Rourke, *Art gallery theorems and algorithms*, Oxford University Press, 1987.
2. J. Urrutia, Art Gallery and illumination problems, Handbook of Computational Geometry (Ed. J.-R. Sack and J. Urrutia), Elsevier Science, pp. 973-1027, 2000.

Approximation algorithms

1. S. K. Ghosh, Approximation algorithm for art gallery problems, Proceedings of the Canadian Information Processing Society Congress, pp. 429-434, 1987. Running time: $O(n^5 \log n)$ time. Approximation ratio: $O(\log n)$. Recently, the running time has been improved to $O(n^4)$ for simple polygons and $O(n^5)$ for polygons with holes.
2. A. Efrat and S. Har-Peled, *Guarding galleries and terrains*, Information Processing Letters, 100 (2006), 238-245. (i) For simple polygons, $O(nc_{opt}^2 \log^4 n)$ expected time and $O(\log c_{opt})$ approx. ratio. (ii) For polygons with h holes, $O(nhc_{opt}^3 polylog\ n)$ expected time and $O(\log n \log(c_{opt} \log n))$ approximation ratio.
3. A. Deshpande, T. Kim, E. D. Demaine1 and S. E. Sarma, *A pseudopolynomial time $O(logn)$-approximation algorithm for art gallery problems*, Proceedings of WADS, LNCS, Springer, no. 4619, pp. 163-174, 2007. Running time: Polynomial in n, the number of walls and the spread, where the spread can be exponential. Approximation ratio: $O(\log c_{opt})$.

Optimal exploration and the Art Gallery Problem

- Suppose an optimal exploration algorithm for a point robot has computed visibility polygons from points p_1, p_2,...., p_k.
- We know that (i) $\bigcup_{i=1}^{k} V(P, p_i) = P$, (ii) $p_i \in V(P, p_j)$ for some $j < i$ and (iii) k is minimum. So, P can be guarded by placing stationary guards at p_1, p_2,...., p_k.
- The exploration problem for a point robot is the Art Gallery problem for stationary guards with additional visibility constraint (ii).
- Our exploration algorithm for a point robot is an approximation algorithm for this variation of the Art Gallery problem.

Open Problem

Can one prove that the exploration problem, like the Art Gallery problem, is NP-hard?

Stationary guards in polygons with additional visibility constraint

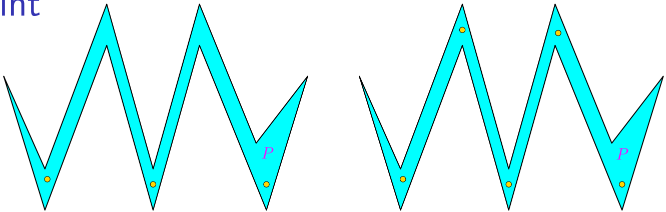

- In the standard art gallery problem, $\lfloor \frac{n}{3} \rfloor$ stationary guards are sufficient and sometime necessary for guarding P containing no holes.
- Suppose, guards g_1, g_2,...., g_k are placed in P for security reasons in a such way that each guard g_i for $i > 1$ is visible at least from one other guard g_j for $i < j$.
- In that case, $\lfloor \frac{n}{3} \rfloor$ guards are not sufficient as $\lfloor \frac{n}{2} \rfloor - 1$ guards are not only necessary but also sufficient.

1. G. Hernandez-Penalver, *Controlling guards*, Proceedings of the Sixth Canadian Conference on Computational Geometry, pp. 387-392, 1994.

- In the standard art gallery problem for a polygon P contains h holes, $\lfloor \frac{n+h}{3} \rfloor$ stationary guards are sufficient and sometime necessary.

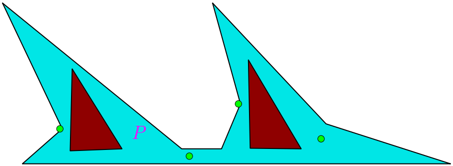

- If the guards also have to satisfy the visibility constraint between them as stated above, then $\lfloor \frac{n+h}{3} \rfloor$ guards are not sufficient.

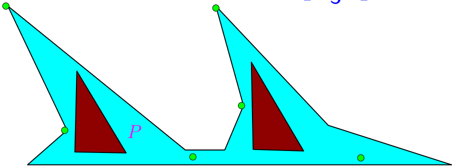

- We conjecture that $\lfloor \frac{n+2h}{3} \rfloor$guards are sufficient for this problem.

1. J. Urrutia, Art Gallery and illumination problems, Handbook of Computational Geometry (Ed. J.-R. Sack and J. Urrutia), Elsevier Science, pp. 973-1027, 2000.

Wathchman route in a polygon

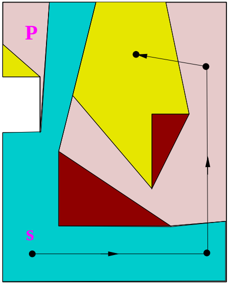

- A *watchman route* in a polygon P is a polygonal path such that every point of P is visible from some point on the path.
- The path of the robot produced by our exploration algorithm is a watchman route inside P.
- This path can also be used as an inspection path of autonomous inspection for subsequent traversal.

1. T. Danner and L. E. Kavraki, *Randomized planning for short inspection paths*, Proceedings of the IEEE International Conference on Robotics and Automation, pp. 971-976, 2000.

Exploring an unknown polygon: Bounded visibility

Computer vision range sensors or algorithms, such as stereo or structured light range finder, can reliably compute the 3D scene locations only up to a depth R. The reliability of depth estimates is inversely related to the distance from the camera. Thus, the range measurements from a vision sensor for objects that are far away are not at all reliable.

Therefore, the portion of the boundary of a polygonal environment within the range distance R is only considered to be visible from the camera of the robot.

Vertices of restricted visibility polygon from p_i with range R are $u_1, u_2, \ldots, u_{12}$.

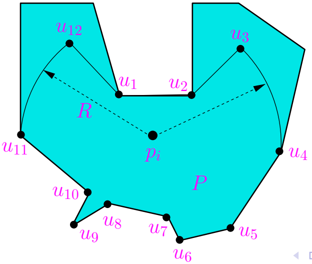

An exploration algorithm using restricted visibility

- The algorithm starts by computing the restricted visibility polygon $RVP(P, p_1)$ from the starting position p_1.

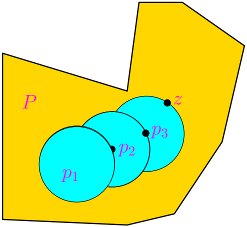

- It chooses the next viewing point p_i on a constructed edge or a circular edge of $RVP(P, p_{i-1})$ for $i \geq 1$ till a boundary point z of P becomes visible.

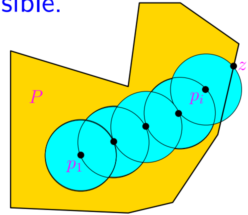

- Taking z as the next viewing point p_i, $RVP(P, p_i)$ is computed. Taking viewing points along the boundary of P in this fashion, restricted visibility polygons are computed till all points of this boundary of P become visible.

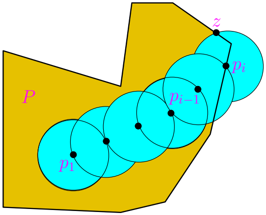

- The process of computing restricted visibility polygons ends once the entire P is explored.

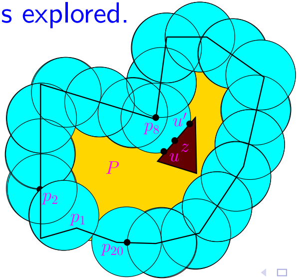

Competitive ratio

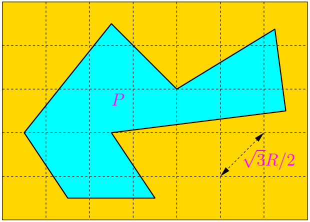

The maximum number of views needed to explore the unknown polygon P with h obstacles of size n is bounded by

$$\left\lfloor \frac{8 \times Area(P)}{3 \times R^2} \right\rfloor + \left\lfloor \frac{Perimeter(P)}{R} \right\rfloor + r + h + 1.$$

The competitive ratio of the algorithm is

$$\left\lfloor \frac{8\pi}{3} + \frac{\pi R \times Perimeter(P)}{Area(P)} + \frac{(r+h+1) \times \pi R^2}{Area(P)} \right\rfloor.$$

Open problem

Can one improve the competitive ratio of the algorithm?

Exploration and Coverage Algorithms

1. A. Bhattacharya, S. K. Ghosh and S. Sarkar, *Exploring an Unknown Polygonal Environment with Bounded Visibility*, Lecture Notes in Computer Science, No. 2073, pp. 640-648, Springer Verlag, 2001.
2. S. K. Ghosh, J. W. Burdick, A. Bhattacharya and S. Sarkar, *On-line algorithms for exploring unknown polygonal environments with discrete visibility*, Special issue on Computational Geometry approaches in Path Planning, IEEE Robotics and Automation Magazine, 2008 vol. 15, no. 2, pp. 67-76, 2008.
3. E. U. Acar and H. Choset, *Sensor-based coverage of unknown environments: Incremental construction of morse decompositions*, The International Journal of Robotics Research, 21 (2002), 345-366.
4. K. Chan and T. W. Lam, *An on-line algorithm for navigating in an unknown environment*, International Journal of Computational Geometry and Applications, 3 (1993), 227-244.

5. H. Choset, *Coverage for robotics– A survey of recent results*, Annals of Mathematics and Artificial Intelligence, 31 (2001), 113-126.
6. X. Deng, T. Kameda and C. Papadimitriou, *How to learn an unknown environment I: The rectilinear case*, Journal of ACM, 45 (1998), 215-245.
7. F. Hoffmann, C. Icking, R. Klein and K. Kriegel, *The polygon exploration problem*, SIAM Journal on Computing, 31 (2001), 577-600.
8. C.J. Taylor and D.J. Kriegman, *Vison-based motion planning and exploration algorithms for mobile robot*, IEEE Transaction on Robotics and Automation, 14 (1998), 417-426.
9. P. Wang, *View planning with combined view and travel cost*, Ph. D. Thesis, Simon Fraser University, Canada, 2007.

Concluding remarks

In this talk, we have presented exploration algorithms under discrete visibility. It can be seen that our algorithms can be implemented easily. We expect that our algorithm will also perform efficiently in practice.