

Scalable visibility color map construction in spatial databases



Farhana Murtaza Choudhury^a, Mohammed Eunus Ali^{a,*}, Sarah Masud^b,
Suman Nath^c, Ishat E. Rabban^a

^a Department of CSE, Bangladesh University of Engineering and Technology, Dhaka, Bangladesh

^b Department of Computer Science, University of Virginia, VA, United States

^c Microsoft Research, Redmond, WA, United States

ARTICLE INFO

Article history:

Received 14 November 2013

Received in revised form

10 December 2013

Accepted 14 December 2013

Recommended by: D. Shasha

Available online 21 December 2013

Keywords:

Spatial databases

Query processing

Three-dimensional (3D) objects

Visibility color map

ABSTRACT

Recent advances in 3D modeling provide us with real 3D datasets to answer queries, such as “What is the best position for a new billboard?” and “Which hotel room has the best view?” in the presence of obstacles. These applications require measuring and differentiating the visibility of an object (target) from different viewpoints in a dataspace, e.g., a billboard may be seen from many points but is readable only from a few points closer to it. In this paper, we formulate the above problem of quantifying the visibility of (from) a target object from (of) the surrounding area with a *visibility color map* (VCM). A VCM is essentially defined as a surface color map of the space, where each viewpoint of the space is assigned a color value that denotes the visibility measure of the target from that viewpoint. Measuring the visibility of a target *even* from a single viewpoint is an expensive operation, as we need to consider factors such as distance, angle, and obstacles between the viewpoint and the target. Hence, a *straightforward approach* to construct the VCM that requires visibility computation for every viewpoint of the surrounding space of the target is prohibitively expensive in terms of both I/Os and computation, especially for a real dataset comprising thousands of obstacles. We propose an efficient approach to compute the VCM based on a key property of the human vision that eliminates the necessity for computing the visibility for a large number of viewpoints of the space. To further reduce the computational overhead, we propose two approximations; namely, *minimum bounding rectangle* and *tangential* approaches with guaranteed error bounds. Our extensive experiments demonstrate the effectiveness and efficiency of our solutions to construct the VCM for real 2D and 3D datasets.

© 2013 Elsevier Ltd. All rights reserved.

1. Introduction

Recent advances in large-scale 3D modeling have enabled capturing urban environments into 3D models. These 3D models give photo-realistic resemblance of urban objects such as buildings, trees, and terrains and are widely used by popular 3D mapping services, e.g.,

Google Maps, Google Earth, and Bing Maps. The increasing availability of these realistic 3D datasets provides us an opportunity to answer many real-life queries involving visibility in the presence of 3D obstacles. For example, an advertising company may wish to determine the visibility of their existing billboards from the surrounding areas in order to find a suitable location to place a new billboard; the police may check the visibility of a surveillance camera to find how well it covers its surrounding areas; and an apartment buyer may want to check the visibility of the nearby sea-beach and mountains from various available apartments.

* Corresponding author. Tel.: +88 029665612; fax: +88 029665612.

E-mail addresses: farhanamc@gmail.com (F.M. Choudhury), eunus@cse.buet.ac.bd (M.E. Ali), sp9cx@virginia.edu (S. Masud), sumann@microsoft.com (S. Nath), ieranik@yahoo.com (I.E. Rabban).

In this paper, we investigate efficient techniques to answer the underlying query required by the above applications: computing visibility of an object (e.g., billboard) from the surrounding continuous space, or that of the surrounding space from a source viewpoint (e.g., camera). Our target applications treat visibility as a continuous notion—e.g., a billboard may be more visible from one location than another, depending on factors such as distance, viewing angle, and obstacles between the viewpoint and the target. Such quantification of visibility is important, because a billboard can be visible from a number of places, but may not be readable from all of them. We therefore use a *visibility function* that provides real-valued visibility measures of various points in the (discretized) 3D space, where the visibility measure of a point denotes its visibility from the viewpoint or to the target object. Thus, the answer to our target query is essentially the visibility measures for every point in the 3D space. The result can be graphically represented as a *heat map*, by assigning colors to various points according to their visibility measures. We call this a *visibility color map* (VCM) of the space for a given target or for a given viewpoint.

Recent works have shown how database techniques can enable efficiently answering various types of visibility queries in the presence of obstacles. Various nearest neighbor (NN) queries consider visibility of objects [1–3]; for example, the visible nearest neighbor query [1] finds the nearest neighbors that are visible from the source. However, these works, like various other computer graphics works [4–9], treat visibility as a binary notion: a point is either visible or not from another point. In contrast, in our target applications, visibility is a continuous notion. Recently, Masud et al. proposed techniques for computing continuous visibility measure of a target object from a particular point in 3D space (e.g., computing visibility of a billboard from a given location) [10]. On the contrary, our target applications require visibility calculation from or of a continuous space, not from a user specified location.

One straightforward way to generate a VCM is to discretize the 3D space and to use the techniques in [10] to compute visibility measure for each discrete point in the space. However, this can be prohibitively expensive. For example, discretizing the surrounding space into 1000 points in each dimension would give a total of 10^9 points in the 3D space; and computing visibility measure for each point by using techniques in [10] would take 128 days! The huge cost comes from two sources: (i) computing the visibility measure based on the distance and angle from all viewpoints, which is computationally expensive and (ii) accessing a large set of obstacles from the database, which is I/O expensive.

We address the above challenges with a three-step solution that uses several novel optimizations to reduce computational and I/O overhead. First, we partition the dataspace into a set of *equi-visible cells*, i.e., all points inside a cell have equal visibility of the target object in terms of visual appearance. We exploit the key observation that when a lens (e.g., a human eye) sees an object without any obstacles, it cannot differentiate between its visual

appearances from a spatially close set of points within an angular resolution (or spatial resolution) of ≈ 4 arcmin ($\approx 0.07^\circ$) [11]. Thus, we can safely prune the visibility computation for a large number of viewpoints within the angular or spatial resolution without affecting viewer's perception. This optimization significantly reduces the computation cost, as we can compute only one visibility measure for each cell.

In the next step, we consider the effect of obstacles. We compute *visible regions*, the regions in the space from where the target object is completely visible in the presence of obstacles. In the final step, we assign visibility measures to these regions from the corresponding cells by spatial joins. Both steps are I/O and computation intensive. For example, they both require retrieving a large number of cells and obstacles from the spatial database. To reduce I/O costs, we employ various indexing techniques to incrementally retrieve a small number of obstacles and cells near the target object. These steps also require performing many computationally expensive operations such as polygon intersections of irregular-shaped regions and cells. To reduce such costs, we represent regions with regular shaped polygons in a quad-tree. We also propose two approximations that further reduce the cost while providing guaranteed small error bounds.

We have evaluated the performance of our solution with real 3D maps of two big cities. We compare our solution with a *baseline* approach that divides the space into a regular shape grid of 500 cells in each dimension and computes visibility from each grid cell. The baseline approach results into more than 30% error while requiring about 800 times more computation time and six orders of magnitude more I/O than our solution. Hence, in the baseline approach, dividing the space into more cells for more accuracy is not feasible for practical applications. On the other hand, our approach provides efficient and effective solution.

In summary, we make the following contributions:

- We formulate the problem of efficiently constructing a visibility color map (VCM) in the presence of obstacles in 2D and 3D spaces.
- We propose an efficient solution to construct a VCM. The solution uses various novel optimizations to significantly reduce the computational and I/O overhead compared to a baseline solution.
- We propose two approximations with guaranteed error bounds and reduced computation to construct the VCM.
- We conduct extensive experiments in two real datasets to show the effectiveness and efficiency of our approaches.

2. Problem formulation and preliminaries

2.1. Problem formulation

The construction of a *visibility color map* (VCM) can be seen from two perspectives: *target-centric* VCM and *viewer-centric* VCM. The construction of a *target-centric*

VCM involves determining how much a given target is visible from every point¹ of the space. On the other hand, a viewer-centric VCM involves determining how much visible each point in the surrounding space is from a given viewer's location.

In both cases, we need to compute visibility and produce a color map of the space where each point of the space is assigned a color value that corresponds to the visibility measure of that point.

Definition 2.1 (*Visibility color map (VCM)*). Given a d -dimensional dataspace \mathcal{R}^d and a set O of obstacles, the VCM is a color map C , where for each point $x \in \mathcal{R}^d$, there exists a color $c_x \in [0, 1]$. The color c_x corresponds to the visibility of a given target object T from x in the case of target-centric VCM and the visibility of x from a given viewer's location q in the case of viewer-centric VCM. Here, c_x is normalized between 0 and 1.

Without loss of generality, we limit our discussion to the construction of target-centric VCM in the subsequent sections. However, our solution is also applicable to the viewer-centric VCM construction, as explained in Section 4.2.

The core of computing a VCM is computing visibility of the target from various points in the space. The most common measure of visibility (or, the perceived size) of an object is the visual angle [12], which is the angle imposed by the viewed object on a lens. The visual angle mainly depends on the characteristics of the viewing lens as well as the distance, angle, and obstacles between the viewer and the target [13].

2.2. Factors affecting visibility

2.2.1. Relative position of the lens and the target

The perceived visibility of a target object mainly depends on the relative position the viewing lens and the target. For a specific target, the visibility varies with the change of distance and angle between the lens and the target.

Distance: If the distance between a target T and a lens increases, the perceived size of T becomes smaller. This is because the visual angle imposed by an object decreases with the increase of the distance between that object and the viewer. As shown in Fig. 1, AB is a target object of length S and the position of a lens is O . When the midpoint of AB is at an orthogonal distance D from O , the visual angle V is calculated using the following formula [14]:

$$V = 2 \arctan\left(\frac{S}{2D}\right) \quad (1)$$

Angle: The perceived size of a target T depends on the angle α between the lens and T [15]. If an object is viewed from an oblique angle, the perceived size of that object becomes smaller than the original size. For the equidistant positions of the lens from T , the visual angle V varies for different values of α . Let AB be a line of length S . A line qm connecting a point q and the midpoint m of AB creates an

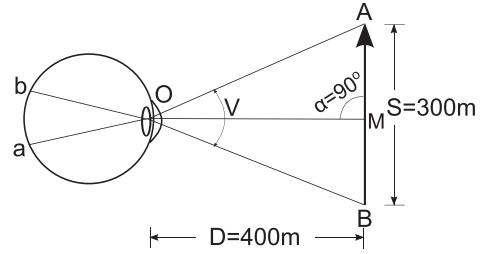


Fig. 1. Effect of distance and angle on visibility.

angle α with AB . If $\alpha = 90^\circ$, the perceived size of AB from a nominal distance is the same as the original length, S . Otherwise, according to the concept of oblique projection [15], the perceived length S_α of AB from q is

$$S_\alpha = \frac{\alpha}{90^\circ} \times S \quad (2)$$

Thus, if we consider the effects of both distance and angle between the target and the lens, the visual angle V can be expressed as $V = 2 \arctan(S_\alpha/2D)$. In Fig. 1, the visibility of the target AB of length S is the visual angle V imposed at the lens O . As in this case $\alpha = 90^\circ$, so the visual angle V is $2 \arctan(300 \text{ m}/(2 \times 400 \text{ m})) = 41.11^\circ$. If $\alpha = 45^\circ$, then V is $2 \arctan(150 \text{ m}/(2 \times 400 \text{ m})) = 21.23^\circ$.

Besides the relative position of the lens, the visibility of a target is also affected by the presence of obstacles.

2.2.2. Obstacles

To show the effect of obstacles on the visibility, we first define the term point to point visibility.

Definition 2.2 (*Point to point visibility*). Given two points p , p' and a set O of obstacles in a space, p and p' are visible to each other if and only if the straight line connecting them, $\overline{pp'}$, does not cut through any obstacle, i.e., $\forall o \in O, \overline{pp'} \cap o = \emptyset$.

Based on the definition of point to point visibility, we formally define the obstructed region as

Definition 2.3 (*Obstructed region*). Given a set O of obstacles, a bounded region R , and a target T , the obstructed region is the set of points where for each point p , (i) p is in R and (ii) p is not point to point visible to all points of T .

The obstructed region contains viewpoints from where the target object is not completely visible. Thus, we only need to measure the visibility for the viewpoints residing outside the obstructed region that form the visible region, and assign colors to these viewpoints of the visible region according to the defined visibility measure (i.e., visual angle).

3. Constructing a VCM

To construct a VCM for a given target T , we need to determine the visibility of T from all discrete points in the surrounding space \mathcal{R} in the presence of a set O of obstacles. We represent visibility of each point $x \in \mathcal{R}$ with a color c_x , which is proportional to the visibility

¹ We have used point and viewpoint interchangeably.

of T from x . We use the terms *visibility measure* and *color* interchangeably.

One naïve approach to compute a VCM is to compute visibility of T from every single viewpoint $x \in \mathcal{R}$. Depending on how finely we discretize the space \mathcal{R} , there can be a large number of points, making the process prohibitively expensive. The high overhead of this naïve approach comes due to the expensive visibility computation from a large number of points and expensive I/O operations to retrieve a large collection of obstacles from a spatial database.

To address these problems, we propose an efficient solution to construct a VCM. The key insights of our solution come from the following two observations. First, human eye cannot visually differentiate a target from viewpoints in close proximity of each other, which eliminates the necessity of computing visibility for all viewpoints in the surrounding space. Second, in most cases, only a small subset of obstacles affect visibility of the target, and thus retrieval of all obstacles is redundant. Such redundancy can be avoided by using various indexing techniques.

In the rest of the section, we describe how we exploit these observations and propose an efficient approach to compute a VCM with reduced computational and I/O overhead. Our approach consists of three steps:

1. First, we partition the space into several *equi-visible cells* in the absence of obstacles exploiting the limitations of human vision. This enables us to compute one single visibility measure for each cell. This significantly reduces computational overhead in contrast to computing the color of each discrete viewpoint (Section 3.1).
2. Second, we compute the effect of obstacles and divide the surrounding space into a set of *visible regions* such that the target is completely visible from a viewpoint if and only if it is within a visible region. To reduce the I/O overhead, we index all obstacles and incrementally retrieve only the potential obstacles that can affect the visibility of the target (Section 3.2).
3. Finally, we join colors computed in the first step (that ignores obstacles) and visible regions computed in the second step to compute a VCM, i.e., colors of different parts of the space in the presence of obstacles. To reduce I/O overhead of retrieving results from the first two steps, we employ various indices (Section 3.3).

For ease of explanation, we assume a 2D space and a target with the shape of line in the subsequent sections. However, our approach is applicable to any target shapes in 2D and 3D spaces. Table 1 lists the notations that we use.

3.1. Partitioning space into equi-visible cells

As mentioned before, the ability of a human eye (or a lens) to distinguish the variation of small details of a target T of size S is limited by its angular resolution μ . To exploit this observation, we partition the space into a set ζ of n *equi-visible cells* $\{\zeta_1, \zeta_2, \dots, \zeta_n\}$. Each cell ζ_i is constructed in

Table 1

Notations used and their meanings.

Notation	Meaning
T	The target object
O	A set of b obstacles $O = \{o_1, o_2, \dots, o_b\}$
V	The visual angle imposed in a lens by T
μ	The angular resolution of a lens
α	The angle between T and the line connecting the lens and the midpoint of T

a way so that the deviation in visibility of T from the viewpoints inside a cell, measured as visual angle V , is not visually perceivable. Hence for any two points $p, p' \in \zeta_i$, visibility from p and p' is perceived as the same if

$$|V_p - V_{p'}|_{p, p' \in \zeta_i} \leq \mu \quad (3)$$

Note that visibility is a *symmetric* measure: visibility of a target at location p from a viewer (i.e., the lens) at location q is the same as that of a target at q from location p . Thus, visibility of a target T from the surrounding space is the same as visibility of the space from T 's location. Therefore, in computing visibility of T from the space, we consider the viewer at the target's location and compute visibility of the space from that location.

Since the value of visual angle depends on the distance and angle between the lens and the target, the partitioning is done in two steps: distance based partition and angle based partition.

3.1.1. Distance based partition

As the perceived size of T varies with the change of distance between T and the viewer's location, our aim is to find a set D of m distances, $D = \{d_0, d_1, \dots, d_{m-1}\}$, where for each pair of points p, p' between d_i and d_{i+1} , $0 \leq i < m-2$, the variation of the perceived visibility from p and p' is less than or equal to the angular resolution μ . Note that, since the visibility varies as α deviates from 90° (as explained in Section 2.2.1), we set $\alpha = 90^\circ$ as the default value for the distance based partitioning.

Partitioning starts from the near point distance d_0 , as a lens cannot focus on any object that is nearer than d_0 [16]. Initially, the visual angle V_0 from the distance d_0 is calculated using Eq. (1). Then, starting from V_0 , the value of the visual angle is decreased by the amount of μ at each step and the corresponding d_i is calculated. When the imposed visual angle from a distance d_{m-1} is less than μ , the distance based partitioning process terminates as for any point farther than d_{m-1} (i.e., d_{max}), the perceived visibility is indistinguishable to the viewing lens. So, we have a set $\{d_0, d_1, \dots, d_{m-1}\}$ of distances where every range $\langle d_i, d_{i+1} \rangle$ is a distance based partition. Fig. 2(a) shows the distance based partitions for a target T , where d_0 is the near point distance and the distance based partitions are $\langle d_0, d_1 \rangle$, $\langle d_1, d_2 \rangle$, and so on.

At this stage, we assign a single color for every distance based partition $\langle d_i, d_{i+1} \rangle$. However, every point in a distance based partition does not perceive the same visibility of the target, e.g., in Fig. 2(a), two viewpoints q_1 and q_2 at a same distance partition may have different perceived visibility due to different viewing angles. Thus, in the next

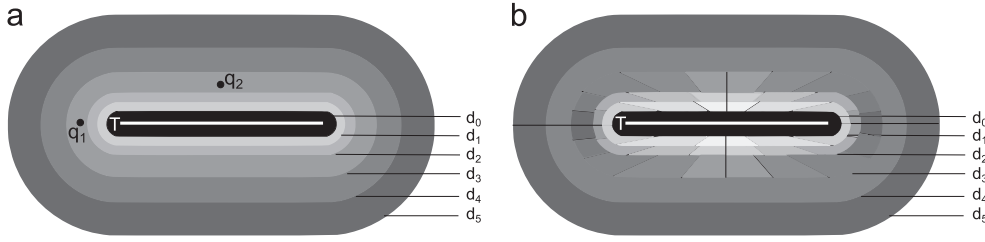


Fig. 2. Partitioning space into equi-visible cells. (a) Partitioning for distance. (b) Partitioning based on distance and angle.

section we incorporate the effect of viewing angle and partition the space into equi-visible cells.

Algorithm 1. Partitioning (T, μ, N_e) .

Input: Target T , angular resolution μ , near point distance N_e .

Output: $\zeta = \{\zeta_1, \zeta_2, \dots, \zeta_n\}$, where ζ_i is a cell resulting from distance and angle based partitioning.

```

1.1 Initialize  $L$  to an empty list;
1.2 Initialize  $\zeta$  to an empty list;
1.3  $S_0 \leftarrow \text{lengthof}(T)$ ;  $i \leftarrow 0$ ;
1.4  $\alpha_0 \leftarrow 90^\circ$ ;  $d_0 \leftarrow N_e$ ;
1.5  $V \leftarrow 2 \arctan \frac{S_0}{2 \times d_0}$ ;
1.6 While  $V \geq \mu$  do
1.7    $d_{i+1} \leftarrow S_0 / (2 \tan(V/2))$ ;
1.8    $L \leftarrow \text{insert}(d_i, d_{i+1})$ ;
1.9    $V \leftarrow V - \mu$ ;  $d_i \leftarrow d_{i+1}$ ;  $i \leftarrow i + 1$ ;
1.10 for each element  $\langle d_i, d_{i+1} \rangle$  of  $L$  do
1.11    $d'_i \leftarrow \text{average}(d_i, d_{i+1})$ ;
1.12    $V_0 \leftarrow 2 \arctan \frac{S}{2 \times d'_i}$ ;
1.13    $\gamma_{i,0} \leftarrow 90^\circ$ ;  $j \leftarrow 1$ ;
1.14    $V \leftarrow V_0 - \mu$ ;
1.15   while  $V \geq \mu$  do
1.16      $S \leftarrow 2 \times d'_i \times \tan(V/2)$ ;
1.17      $\alpha \leftarrow (S \times 90^\circ) / S_0$ ;
1.18      $\gamma_{ij} \leftarrow \gamma_{ij-1} - \alpha$ ;
1.19      $\zeta \leftarrow \text{insert}(\langle d_i, d_{i+1}, \gamma_{ij-1}, \gamma_{ij} \rangle)$ ;
1.20      $V \leftarrow V - \mu$ ;  $j \leftarrow j + 1$ ;
1.21 return  $\zeta$ ;
```

3.1.2. Angle based partition

For each distance based partition $\langle d_i, d_{i+1} \rangle$, $0 \leq i < m-2$, we get l_i numbers of angle based partitions $\langle \gamma_{ij}, \gamma_{ij+1} \rangle$, $0 \leq j < l_i-2$, where the value of l_i is different for each distance based partition. The visibility of every point of a partition $\langle d_i, d_{i+1}, \gamma_{ij}, \gamma_{ij+1} \rangle$ is considered as the same. We call such a partition an *equi-visible cell* (or, just *cell* in short).

As the change in perceived length due to change in α is symmetric with respect to both the parallel and normal axes of the object plane, we compute the angle based partitions only for the first quadrant. The partitions in other quadrants are then obtained by reflecting the partitions of the first quadrant. The procedure of angle based partitioning is as follows:

- (i) For each distance based partition $\langle d_i, d_{i+1} \rangle$, the angle based partitioning starts by initializing $\alpha_0 = 90^\circ$, $\gamma_{i,0} = 90^\circ$, and $S_0 = S$, where the size of T is S . The visual angle V_0 is calculated for α_0 using Eq. (1) where

$D = d'_i$. Here d'_i is the average value of distances d_i and d_{i+1} , i.e., $d'_i = (d_i + d_{i+1})/2$.

- (ii) At each step j of the angle based partitioning for $\langle d_i, d_{i+1} \rangle$, the perceived size S_j is calculated for visual angle $V_j = V_{j-1} - \mu$ and distance d'_i using Eq. (1), where $j \geq 1$, e.g., $S_1 = 2D \times \tan(V_1/2)$, for $j=1$. The angle α_j , for which S_j is perceived, is computed using Eq. (2), e.g., $\alpha_1 = (S_1 \times 90^\circ)/S_0$, for $j=1$. The visual angle V_j is obtained for the change in γ_{ij-1} by the amount of α_j , so γ_{ij} is updated as $\gamma_{ij-1} - \alpha_j$. Thus we get an angle based partition $\langle \gamma_{ij-1}, \gamma_{ij} \rangle$ for $\langle d_i, d_{i+1} \rangle$ at each step.
- (iii) When the perceived visual angle $V_j < \mu$, the angle based partitioning process for a distance based partition terminates.

The above process is repeated for each distance based partition and finally we get the set $\zeta = \{\zeta_1, \zeta_2, \dots, \zeta_n\}$ of n cells where $n = m \times l_i$, $0 \leq i < m-1$. Fig. 2(b) shows the angle-based partitions for three distance-based partitions $\langle d_0, d_1 \rangle$, $\langle d_1, d_2 \rangle$, and $\langle d_2, d_3 \rangle$ for a target object T .

The steps of the distance based and angle based partitioning are shown in Algorithm 1. Lines 1.6–1.9 in Algorithm 1 show the distance based partitioning steps. Each partition from the distance based partitioning is further partitioned for angles. Lines 1.10–1.20 show the angle based partitioning steps.

After both partitioning steps are done, we compute visibility (i.e., color) of each cell. Since all viewpoints within a cell have the same visibility, we assign the color of the entire cell as the visibility of the center of the cell to the target T .

Note that we have not considered the effect of obstacles yet. A caveat of this is that in the presence of obstacles, the target T may not be visible from an entire cell or parts of a cell even if the cell is assigned a good visibility value. We address this next by considering the effect of obstacles.

3.2. Computing the effects of obstacles

Given a target object T and a set O of obstacles, we would like to determine the set S of *visible regions* surrounding T such that T is completely visible from a viewpoint q if and only if q is inside a visible region.

A naïve approach to determine the effects of obstacles is to retrieve all obstacles of O and calculate the corresponding changes in the visibility. But this approach is prohibitively costly in terms of both I/O and computation, especially in the presence of a large number of obstacles.

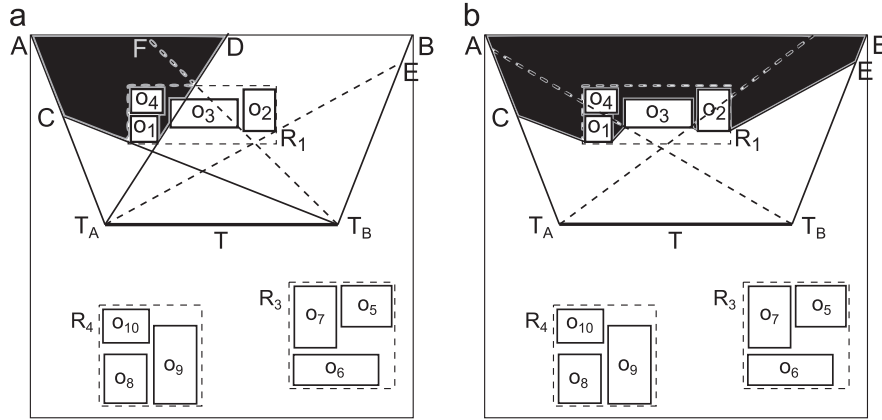


Fig. 3. Visible region construction.

Moreover, considering all obstacles in the database is wasteful as only a relatively small number of obstacles around T affect visibility. To efficiently retrieve this small number of obstacles around T , we index all obstacles in an R^* -tree [17], a variation of R -tree [18]. An R^* -tree consists of a hierarchy of *Minimum Bounding Rectangles* (MBRs), where each MBR corresponds to a tree node and bounds all the MBRs in its sub-tree. Data objects (obstacles, in our case) are stored in leaf nodes.

A recent work [10] has developed a technique to determine the obstacles that affect the visibility of the target from a specific viewpoint. On the contrary, we need to compute the visible region of the whole space instead of a specific viewpoint with respect to a target. Thus, we cannot adopt the computation of the visible region from [10]. Our approach to compute the visible region in the presence of large number of obstacles is as follows.

Initially the set V_R of visible regions contains the region of R that is covered only by the field of view (FOV) with respect to the target T . As shown in Fig. 3(a), V_R is the region bounded by the points T_A , T_B , A , and B where T_A and T_B are the corner points of T . Here, $FOV = 120^\circ$, the usual FOV of the human eye [19]. Initially, the set O_R of obstructed regions is the region of R that is outside the FOV. As T cannot be viewed from the region outside the FOV, the obstacles residing in this region is discarded from consideration.

Next, we refine the set V_R and the set O_R by considering one obstacle at a time. The obstacle retrieval starts from the root node of the R^* -tree. Only the nodes that intersect with a region in V_R are incrementally retrieved from the R^* -tree according to their non-decreasing distances from the target T . If the retrieved node is an MBR, its elements are further discovered from the R^* -tree. For example, in Fig. 3(a), when the MBR R_1 is accessed, its elements o_1 , o_2 , o_3 , and o_4 are further discovered. If the retrieved node is an obstacle o , the regions in V_R and O_R are updated according to the effect of o . We term the effect of a single obstacle o on visibility as the shadow of o , W_o .

Definition 3.1. Shadow of an obstacle o , W_o . W_o is the region formed by the set of points P , where for any point p

of P , there is at least a point t_p on T such that the line segment joining t_p and p either intersects or touches o .

From each point p of the W_o , T is either completely or partially obstructed. The boundary of a shadow W_o contains exactly two straight lines, which are tangents between the obstacle and the target. If these lines are rays, not the line segments that meet each other, then the region W_o is unbounded. If W_o is unbounded, we consider only the portion that is bounded by the given region R . In Fig. 3(a), the shadow of o_1 is W_{o_1} (shown with black shade), the region bounded by o_1 and the points A , C , and D .

While updating the V_R and O_R for the shadow of a retrieved obstacle, there are three cases to be considered:

- (i) If the obstacle o or its shadow W_o does not overlap with any obstructed region of O_R , we exclude W_o from V_R and include it in O_R . In Fig. 3(a), obstacle o_1 is the first obstacle retrieved according to its non-decreasing distance from T . As there is no other obstructed region to be overlapped with o_1 or its shadow W_{o_1} (shown as the black region), W_{o_1} is now excluded from V_R and included in O_R .
- (ii) If o or W_o overlaps with one or more obstructed regions of O_R , we combine these regions and W_o into a single obstructed region and discard this region from V_R . Let N be the set of W_o and the obstructed regions that overlap with W_o . To combine the regions in N , we determine the leftmost tangent line l and the rightmost tangent line r of all the shadows of N such that the region bounded by l , r and the union of the shadows of N enclose all the obstructed regions of N . This shadow resembles the combined effect of the obstacles that are included in N . We replace the regions of N from O_R with this combined region and discard it from V_R . As an example, in Fig. 3(a), the shadow W_{o_2} (shown with dotted lines) of the next retrieved obstacle o_2 overlaps with the existing obstructed region (shown with black shade). Here, the leftmost tangent line and the rightmost tangent line of these obstructed regions are the line connecting o_1 , C and the line connecting o_2 , E , respectively.

The region enclosed by these two lines and the union of the shadows is discarded from V_R . Similarly, the next retrieved obstacle o_3 overlaps with O_R . The shadows combined for obstacle o_1 , o_2 , and o_3 are shown in black shade in Fig. 3(b), where the shadow of o_3 is shown with dotted lines.

- (iii) If o is entirely inside any obstructed region, it will not contribute to the visibility. So we discard o from consideration. In Fig. 3(b), the effect of obstacle o_4 is not calculated as it is entirely inside the current obstructed region.

Note that, the visible region includes viewpoints from where the target is entirely visible. This approach is suitable for applications like placement of billboards, where partial visibility of the target from a viewpoint does not make sense. There can be some applications that require finding the viewpoints from where a target is partially visible. We leave computing the partial visible viewpoints as the scope of the future study.

3.3. Merging cells with visible regions

In the final step of producing a VCM, we combine colors computed in the first step (that ignores obstacles) and visible regions computed in the second step to compute color of each cell in the presence of obstacles. Intuitively, all obstructed regions are assigned color zero (representing zero visibility), while all visible regions are assigned colors from their corresponding cells. This requires taking intersection (spatial join) of all cells and visible regions.

The above process can be expensive for two reasons. First, visible regions are usually irregular polygons and intersecting them with cells incurs high computational overhead due to the complex shapes of the polygons. We address this by storing each polygon as a set of regular shapes (rectangles) with a quad-tree [20]. Second, the number of cells can be quite large and intersecting visible regions with all cells can be expensive. We address this by indexing the cells in an R*-tree and intersecting each visible region only with its overlapping cells. We describe these optimizations below.

3.3.1. Indexing visible regions

We index visible (and obstructed) regions in a 2D (3D) space with a quad-tree (octree). Quad-trees (Octrees) partition a 2D (3D) space by recursively subdividing it into four quadrants (eight octants) or blocks. Initially the whole space is represented with a single quad-tree block. In the visible region computation phase, when a region is obstructed due to a retrieved obstacle, a quad-tree block is partitioned into four equal blocks if it intersects with the obstructed region. The partitioning continues until (i) a quad-tree block is completely visible, or (ii) completely obstructed, or (iii) the size of a block is below a threshold value ϑ . The threshold size is determined in the partitioning phase. If the set of the boundary points of an equivalence cell ζ_i is β_{ζ_i} , the minimum distance between any two opposite boundaries over all cells is specified as ϑ , i.e.,

$$\vartheta = \arg \min_{\zeta_i, 1 \leq i \leq n} \left(\arg \min_{p, p' \in \beta_{\zeta_i}} (\text{mindist}(p, p')) \right)$$

where points $p, p' \in \beta_{\zeta_i}$ are points of opposite boundaries.

As ϑ is the minimum size of a cell obtained in the partitioning phase and the deviation in the perceived visibility for a block-size less than ϑ is not distinguishable due to angular resolution of a lens, dividing a quad-tree block into size smaller than ϑ for more accuracy is redundant.

Fig. 4(a) shows the quad-tree of the obstructed region (black blocks) and visible region (white blocks inside FOV). After the quad-tree is constructed, we start assigning color to each of its blocks. We get the VCM after all the blocks are assigned colors. The blocks that fall in the obstructed region are assigned zero visibility. The remaining blocks (i.e., visible region) are assigned colors based on the colors of their corresponding cells, as described next.

3.3.2. Indexing cells and constructing VCM

For each quad-tree block with unassigned color, we need to find all overlapping cells in order to find its color. To expedite this process, we index cells with an R*-tree that we call the *color-tree*. Leaf nodes in the color-tree represent cells and non-leaf nodes represent MBRs containing the shapes of their children nodes. Then, for each

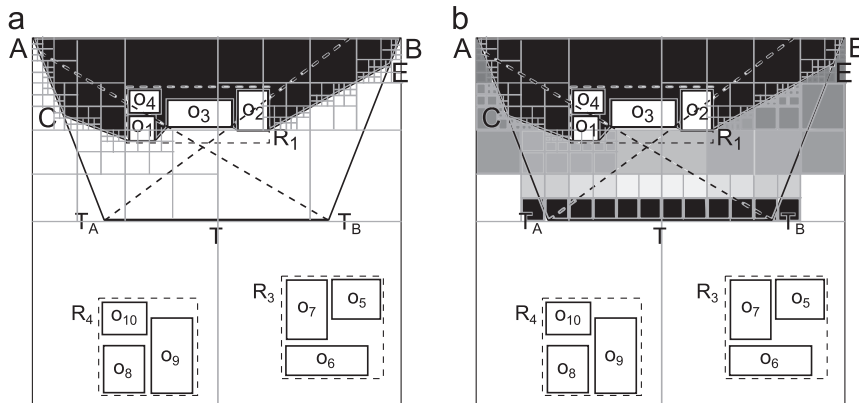


Fig. 4. Construction of visibility color map.

quad-tree block, we run a range query on the color-tree. The color for an unassigned quad-tree block is obtained by calculating the intersection of the spatial region of that block and the cells from the color-tree. If the quad-tree block intersects with a single cell of the color partition, the block is assigned the color of that cell. If a quad-tree block intersects with multiple cells, in a 2D (3D) space we further divide that block into four (eight) equal blocks. The division is continued until either a block intersects with a single cell or the size of a block is below the threshold value ϑ (Section 3.3.1). The process terminates when all quad-tree blocks of the visible region are colored according to the visibility measure. Fig. 4(b) shows the resulting VCM constructed by combining the color partitions of Fig. 2(b) and the quad-tree of Fig. 4(a).

The steps of constructing a VCM are shown in Algorithm 2. Lines 2.5–2.8 show accessing the color-tree nodes. If an accessed node is an MBR, its elements are further discovered from the color-tree (lines 2.20–2.21). If the accessed node is a leaf node, the quad-tree nodes that are not colored yet and intersect with this node are either colored or partitioned further (lines 2.14–2.19). Finally, the colored quad-tree is returned as the complete VCM.

Algorithm 2. ConstructVCM($T, FOV, \vartheta, Qtree, CTree$).

```

2.1 Initialize Q to an empty queue;
2.2 Initialize L to an empty list;
2.3 node ← CTree.root; end ← false;
2.4 L ← Get_unassigned_quad_leaf(Qtree, FOV);
2.5 while node ≠ empty and end = false do
2.6   for each element  $n_e$  of node do
2.7     if Inside_visibility_region( $n_e$ ) = true then
2.8       | Enqueue(Q,  $n_e$ );
2.9   continue ← true;
2.10 while continue = true do
2.11   continue ← false; node ← Dequeue(Q);
2.12   if node = empty or L = empty then
2.13     | end ← true;
2.14   else if node is a data object then
2.15     | Initialize Still_unassigned_quads to an empty list;
2.16     for each element  $L_i$  of L do
2.17       | Still_unassigned_quads ← insert(Divide_and_color( $L_i$ , node,  $\vartheta$ ));
2.18     L ← Still_unassigned_quads;
2.19     continue ← true;
2.20   else
2.21     | node ← child(node);
2.22 return Qtree;
```

Algorithm 3. Divide_for_color ($Qblock, Cnode, \vartheta$).

```

3.1 Initialize L to an empty list;
3.2 if intersects (Qblock, Cnode) = true then
3.3   if Totally_inside(Qblock, Cnode) = true or size(Cnode) <  $\vartheta$  then
3.4     | color(Qblock) ← color(Cnode);
3.5   else
3.6     | Divide(Qblock);
3.7     for each children  $q_c$  of Qblock do
3.8       | Divide_for_color( $q_c$ , node,  $\vartheta$ );
3.9   else
3.10    | L ← Qblock;
3.11 return L;
```

Since the color partitioning process results in complex shaped cells (e.g., curves), it is computationally expensive

to combine these shapes with the quad-tree blocks of the visible region. Thus, we propose two approaches to approximate color partition cells.

4. Extensions

In this section, we discuss a few extensions to our basic algorithm described in the previous section.

4.1. Approximation of partitions

The algorithm in the previous section partitions the space according to the relative distance and angle between the lens and the target. Based on these parameters, the cells are bounded by arcs and straight lines. To construct the VCM, we need to compute the intersection of these cells with the quad-tree blocks. The process is computationally expensive due to the complex shape of the cells and the target.

To address this, we introduce two approximations that reduce the computational overhead at the cost of small bounded errors: (i) minimum bounding rectangle (MBR) of a cell and (ii) tangents of the arcs of a cell. For the ease of explanation we analyze errors for targets with regular shapes, e.g., lines without loss of generality. For such targets, the complex shaped partitions are bounded by two arcs and two straight lines. Here we discuss the approximations of color partitions and the maximum error resulting from these approximations.

4.1.1. MBR approximation

An approach to approximate the curves of a cell is to enclose the cell using its MBR where the area covered by the cell is approximated by the area covered by the enclosing MBR. This is illustrated in Fig. 5(a), where a cell consists of two concentric arcs c_i and c_{i+1} (centered at a corner O of the target) of radius r_i and r_{i+1} , respectively. $ABCD$ is the MBR of this cell. Let θ_i and θ_{i+1} be the angles (in degree) created by c_i and c_{i+1} at the center respectively. We denote the area of the MBR as A_r . So the area bounded by the cell, A_p is $(\theta_{i+1}/360 \times \pi r_{i+1}^2) - (\theta_i/360 \times \pi r_i^2)$. Hence the area that gets wrong color due to this MBR approximation is $A_r - A_p$ (shaded region in Fig. 5(a)).

Error bound analysis: For targets with regular shapes such as lines (Fig. 2(b)), the largest cell that can yield the maximum possible error consists of two half circles centered at a corner point of the target. Hence we formulate the maximum error bound by referring to Fig. 5(b). Here, we want to approximate the area of a cell p_i bounded by two half circles c_i and c_{i+1} centered at O , a corner point of target T of length S . Here, c_i and c_{i+1} belong to distance based partition d_i and d_{i+1} , respectively. We approximate p_i by taking its MBR, $ABCD$. So, the darkly shaded region is wrongly colored for MBR approximation. The lightly shaded region is wrongly colored too, but it is considered for distance based partition $\langle d_{i-1}, d_i \rangle$ i.e., for cell p_{i-1} . As the total error is calculated incrementally for each cell, so the error for this lightly shaded region is calculated only once for the cell p_{i-1} . Let, r_i and r_{i+1} be the radii of c_i and c_{i+1} , respectively. Hence, the width of cell p_i

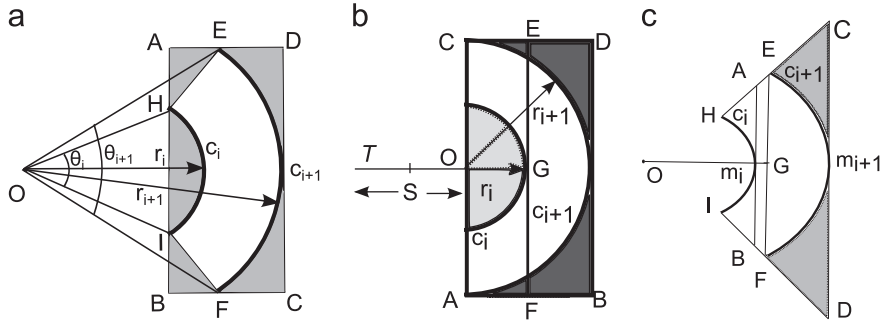


Fig. 5. Approximating a cell using MBR and tangents.

is $r_{i+1} - r_i$. The farthest distance from O to any point of the MBR is $\Delta r = \sqrt{2} \times r_{i+1}$. So the points of the MBR that fall within the distance r_{i+1} and $\sqrt{2} \times r_{i+1}$ from point O get wrongly colored with the p_i 's color. From Eq. (1) we get that the width of the distance based partitions increases with the increase of the distance between the target and the partition. So the maximum number of distance based partitions in d_{i+1} that can be wrongly colored is $n = (\Delta r - r_{i+1}) / (r_{i+1} - r_i)$.

For each distance based partition, corresponding angle based partitions are calculated by taking the angle from the midpoint of the target. So only the angle based partitions that fall within the range $(\arctan r_{i+1}/S/2, -\arctan r_{i+1}/S/2)$ can lie in the darkly shaded region. Let there are a_i such angle based partitions in total. As the difference in visual angle in consecutive partitions is μ , the maximum variation in color for a cell is $\Delta c_{m_i} = n \times a_i \times \mu$. So when the total number of partitions is k and the area of i th cell is A_i , the total error in coloring is

$$E_{MBR} = \sum_{i=1}^k \Delta c_{m_i} \times A_i \quad (4)$$

4.1.2. Tangential approximation

We obtain another approximation approach by taking tangents in the midpoints of the two arcs that encloses a cell. Here the area enclosed by two arcs is approximated by the area enclosed by the tangents at their midpoints. In Fig. 5(c) the cell is bounded by arcs c_i and c_{i+1} centered at a corner O of the target T . We take two tangents AB and CD of c_i and c_{i+1} at their midpoints m_i and m_{i+1} , respectively. So we want to approximate the area bounded by these two arcs with the trapezoid $ABDC$. Let the angles created at O by c_i and c_{i+1} are θ_i and θ_{i+1} (in degrees). If we color the trapezoid $ABDC$ instead of the cell, the region that gets wrong color is the shaded region of Fig. 5(c). Let the area of the trapezoid $ABDC$ be A_{t_1} , the radius of c_{i+1} is r_{i+1} , the midpoint of \overline{EF} is G , and the length of \overline{OG} is x . Then the length of \overline{EG} is $\sqrt{r_{i+1}^2 - x^2}$. The area of the segment bounded by \overline{EG} and c_{i+1} , A_c is $(\theta_{i+1}/360 \times \pi r_{i+1}^2) - x \times \sqrt{r_{i+1}^2 - x^2}$. If the area of the trapezoid $ABFE$ is A_{t_2} , then the area of the shaded region A_e is $A_{t_1} - A_c - A_{t_2}$. According to this tangential approximation the region that is wrongly colored due to cell $HIFE$ is this shaded region with area A_e .

Error bound analysis: In the case of MBR approximation the cell bounded by two half circles is approximated by an enclosing MBR, while for tangential approximation that cell is approximated by a rectangle bounded by the tangents of those two half circles. Referring to Fig. 5(b), the cell bounded by arcs c_i and c_{i+1} is approximated by rectangle $ABDC$ and rectangle $FBDE$ for MBR and tangential approaches, respectively. In this figure, G is the midpoint of c_i and the farthest distance from G to any point inside the approximated cell is $\Delta r = \sqrt{r_{i+1}^2 + (r_{i+1} - r_i)^2}$. So the points inside the darkly shaded region are wrongly colored. As the width of the distance based partitions relates inversely to the distance between the target and the partition (Eq. (1)), the maximum number of distance based partitions in d_{i+1} that can be wrongly colored is $n = (\Delta r - r_{i+1}) / (r_{i+1} - r_i)$. In the case of angle based partitioning, only the angle based partitions that fall within the range $(\arctan r_{i+1}/(S/2 + r_i), -\arctan r_{i+1}/(S/2 + r_i))$ can lie in the darkly shaded region inside the rectangle $EFBD$. Let there are a_i such angle based partitions in total. As the difference in visual angle in consecutive partitions is μ , the maximum variation in color for a cell is $\Delta c_{t_i} = n \times a_i \times \mu$. So when the total number of partitions is k and the area of i th cell is A_i , the total error is

$$E_{Tangent} = \sum_{i=1}^k \Delta c_{t_i} \times A_i \quad (5)$$

In the above sections, we discussed our approach to construct *target-centric visibility color map*. The approach is the same for the *viewer-centric visibility color map* with an additional initialization step. The details of *viewer-centric VCM* is discussed below.

4.2. Viewer-centric VCM

A *viewer-centric VCM* is constructed by calculating the visibility of the surrounding space for a given viewpoint q and a set O of obstacles. Unlike the *target-centric VCM* where a particular target is specified, in the case of the *viewer-centric VCM*, a particular viewer position is specified. To measure the visibility using Eqs. (1) and (2) in the case of the viewer-centric VCM, first we need an initial value of the size S of a target from the given information.

The visibility of any point farther than a distance d_{max} is not visually distinguishable if the perceived visual angle of that point is less than the angular resolution μ from q (Section

3.1.1). Based on this fact, we assume a circular region R of radius d_{max} centered at q to construct the VCM of R only, as the visibility of the points outside R is not distinguishable by a viewer at q . Using Eq. (1), we calculate the size S of a target for which visual angle $V = \mu$ is perceived at distance d_{max} . After taking S as the size of the target, the steps of our approach discussed for the *target-centric* VCM is applied in the same manner to construct the *viewer-centric* VCM.

4.3. Incremental processing

During the computation of the VCM, we have assumed that the field of view (FOV), i.e., the orientation of the viewpoint or the target is fixed. At a particular orientation or gaze direction, only the extents of the space that is inside the FOV are visible. With the change in viewing direction, some areas that were previously outside the FOV become visible. In this case, we do not need to compute the full VCM each time, rather we can incrementally construct the VCM by computing for the newly visible parts only.

In the incremental process, the only information that varies is the gaze direction. So, as a preprocessing step we can construct the color-tree and the visible region quad-tree in the above discussed method by considering the FOV as 360° . A VCM is then constructed by combining the color-tree and the quad-tree for a particular gaze direction. When the gaze direction changes, only the uncolored quad-tree blocks that are now included in the visibility region are assigned colors. Hence the color-tree and the quad-tree are constructed only once. This reduces the computational complexity to a great extent by avoiding same calculations repetitively. From our conducted experiments we also observe that the processing time required to combine the color-tree and the quad-tree is much smaller than the processing time required to construct the visibility region quad-tree for a dataset of densely distributed large number of obstacles (Sections 5.3.1 and 5.4.1). So for such cases we can significantly improve the performance of our proposed solution by adopting the preprocessing strategy.

5. Experimental evaluation

We evaluate the performance of our proposed algorithm for constructing the *visibility color map* (VCM) with two real datasets. Specifically, at first we compare our approach with a *baseline approach* that approximates the total space into a regular grid and compute visibility from the midpoint of each grid cell, and then we compare our two approximation algorithms, i.e., *MBR approximation* (VCM_M) and *tangential approximation* (VCM_T) with the exact method (VCM_E). The algorithms are implemented in C++ and the experiments are conducted on a core i5 2.40 GHz PC with 3 GB RAM, running Microsoft Windows 7.

5.1. Experimental setup

Our experiments are based on two real datasets: (1) British² representing 5985 data objects obtained from

Table 2

Parameters.

Parameter	Range	Default
Angular resolution (μ)	1, 2, 4, 8, 16	4
Minimum block size (θ)	1, 2, 4, 8, 16	1
Query space area (A_Q)	0.05, 0.10, 0.15, 0.2, 0.25	0.15
Field of view (FOV)	60, 120, 180, 240, 300, 360	120
Length of target (L_T)	0.05, 0.10, 0.15, 0.2, 0.25	0.15
Dataset size (D_S)	5k, 10k, 15k, 20k, 25k	

Table 3

Comparison with the baseline.

Dataset	Baseline		VCM_E
	Total time (h)	Error (%)	Total time (s)
British	58	30.33	253.07
Boston	61	31.76	245.26

British ordnance survey³ and (2) Boston⁴ representing 130 043 data objects in Boston downtown. In both datasets, objects are represented as 3D rectangles that are used as obstacles in our experiments. For both datasets, we normalize the dataspace into a span of $10\,000 \times 10\,000$ square units. For 2D, the datasets are normalized by considering the z-axis value as 0. All obstacles are indexed by an R*-tree, with the disk page size fixed at 1 kB.

The experiments investigate the performance of the proposed solutions by varying five parameters: (i) angular resolution (μ) in arcminutes, (ii) the threshold of the quad-tree block size (θ) as the multiple of the calculated minimum size (as explained in Section 3.3.1), (iii) the area of the space (A_Q) as the percentage of the total area, (iv) field of view (FOV) in degrees, and (v) the length of the target (L_T) as the percentage of the length of total dataspace. We have also varied the dataset size using both *Uniform* and *Zipf* distribution of the obstacles. The range and default value of each parameter are listed in Table 2. In concordance with human vision, the default values of μ and FOV are set as 4 arcmin [11] and 120° [19], respectively. The minimum threshold of quad-tree block size as calculated in Section 3.3.1 is used as the default value of θ . The default values of other parameters are set to their median values.

The performance metrics that are used in our experiments are (i) the total processing time, (ii) the total I/O cost, and (iii) the error introduced by the two approximations: VCM_M and VCM_T . We calculate the approximation error as the deviation from the color map of VCM_E , i.e.,

$$error = \frac{\sum_i (c_{e_i} \times A_i - c_{a_i} \times A_i)}{\sum_i c_{e_i} \times A_i} \quad (6)$$

here c_{e_i} is the color of i th cell in VCM_E , c_{a_i} is the color of i th cell in VCM_M or VCM_T , and A_i is the area of the i th cell.

³ <http://www.ordnancesurvey.co.uk/oswebsite/indexA.html>

⁴ http://www.bostonredevelopmentauthority.org/BRA_3D_Models/3D-download.html

² <http://www.citygml.org/index.php?id=1539>

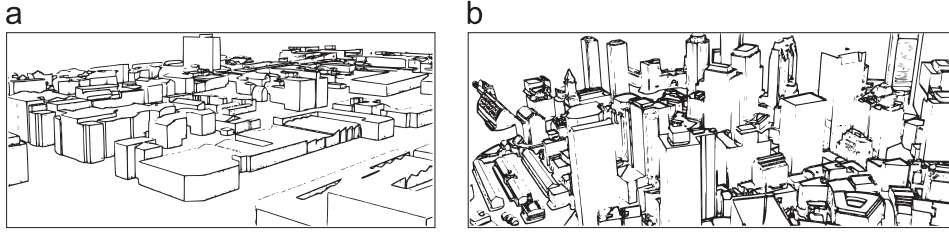


Fig. 6. Real datasets used in the experiments. (a) British ordnance survey. (b) Boston downtown.

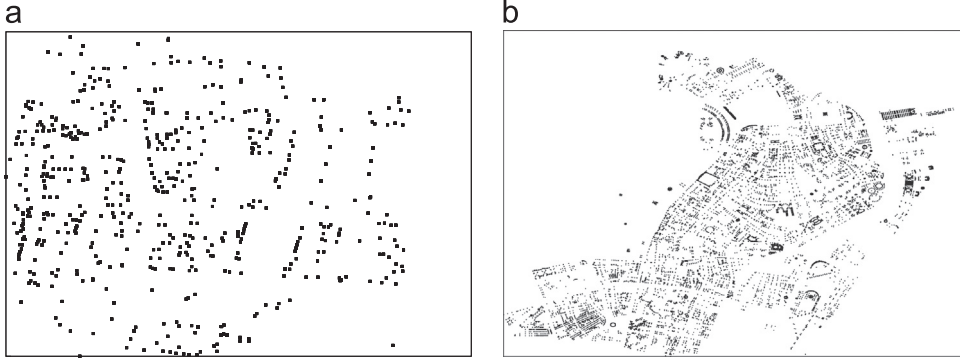


Fig. 7. Dataset distribution. (a) British ordnance survey. (b) Boston downtown.

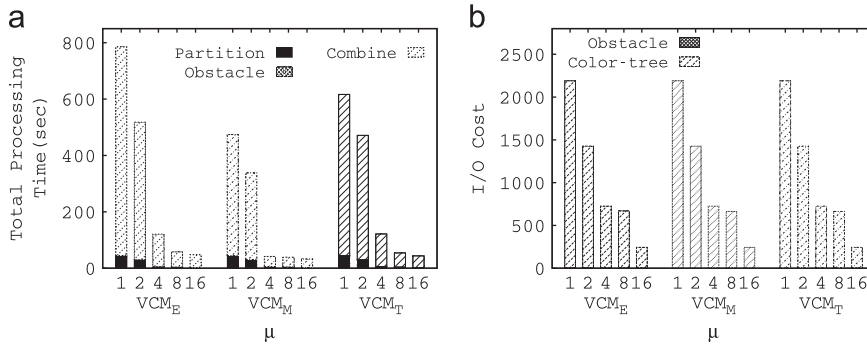


Fig. 8. Effect of μ in 2D British (a and b) dataset.

For each experiment, we have evaluated our solution for the target at 100 random positions and reported their average performance. We have conducted extensive experiments using two datasets for both 2D and 3D spaces.⁵ Since 2D dataspace is a subset of 3D dataspace and most of the real applications involve 3D scenario, we omit the detailed results of 2D.

5.2. Comparison with the baseline

The performance improvement of our approach over the *baseline* approach is measured in terms of the total processing time required for computing the VCM. In the naïve approach, we need to compute visibility of the target from *infinite* number of points to construct the VCM as

every point in the dataspace acts as a viewpoint. Even discretizing the surrounding space into 1000 points in each dimension would give a total of 10^9 points in the 3D space. We observe that the time required for such computation using British and Boston datasets are approximately 74 days and 128 days, respectively. As such naïve approach is trivial, we can further improve it by dividing the dataspace into $500 \times 500 \times 500$ cubic cells in 3D and choose the visibility from the middle point of each cell to *represent* the visibility of that entire cell. In the rest of the paper we refer to this approach as the *baseline* approach to construct the VCM. The comparison between the baseline approach and our proposed solution VCM_E is presented in Table 3 in terms of total processing time and introduced error. In this case the parameters are set to their default values (Table 2). We observe that when compared to VCM_E, for British dataset the baseline approach runs 818 times slower and introduces 30% error, while for Boston dataset the baseline approach runs 837 times slower and introduces 32% error (Figs. 6 and 7).

⁵ Datasets and query sets used in the experiments are available for download from the following link: <http://www.buet.ac.bd/cse/users/faculty/eunus/Datasets.zip>.

Table 4
Error in VCM_M and VCM_T for varying μ in 2D.

Dataset	Method	Error (%)				
		1	2	4	8	16
British	VCM_M	2.0	1.64	4.25	4.42	5.38
	VCM_T	0.01	0.01	0.01	3.64	3.91

Table 5
Errors in VCM_M and VCM_T for ϑ in 2D.

Dataset	Method	Error (%)				
		1	2	4	8	16
British	VCM_M	0.01	2.74	9.29	12.59	12.83
	VCM_T	0.01	0.77	4.37	6.26	7.54

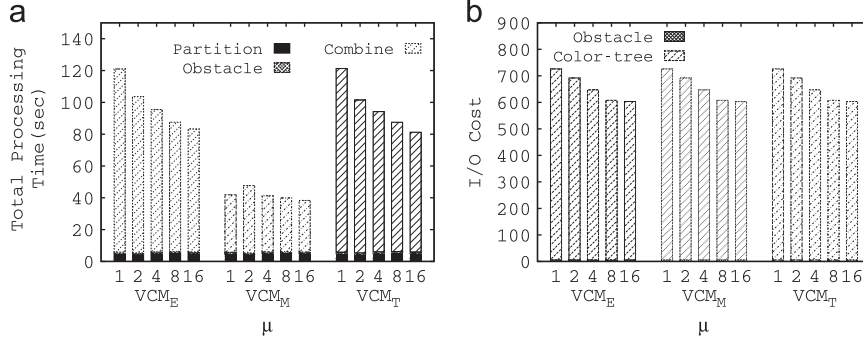


Fig. 9. Effect of ϑ in 2D British (a and b).

5.3. Performance in 2D

In this section, we present experimental results for 2D datasets.

5.3.1. Effect of μ

In this experiment, we vary the value of μ as 1, 2, 4, 8, and 16 arcmin and measure the total processing time and I/Os for British datasets (Fig. 8). We also present the errors resulted in the two approximations (Table 4).

For British dataset, on average VCM_M and VCM_T are 65% times and 17% times faster than VCM_E , respectively. In general, with the increase in μ , the cell size increases and the number of total partition in the dataspace decreases. Moreover, larger μ yields fewer branching in the visibility region quad-tree. So, with the increase in μ , total processing time and I/O cost decrease rapidly for both datasets. The total I/O cost is composed of (i) cost of partitioning the total dataspace to form the color-tree, (ii) cost of retrieving obstacles to form visibility region quad-tree, and (iii) cost of combining the color-tree and quad-tree to form the VCM. As these three costs hardly differ for all three approaches, they result in similar I/O cost. On the other hand, as with the increase in μ , the area of partition cells (A_i) gets larger, the estimated error increases for both VCM_M and VCM_T (Eq. (6)). The average errors introduced in VCM_M and VCM_T are 3.5% and 1.5%, respectively.

5.3.2. Effect of ϑ

In this experiment, we vary the quad-tree block size (ϑ) as 1, 2, 4, 8, and 16 times of the minimum threshold of a quad-tree block. The results are presented in Fig. 9 and Table 5.

For British dataset, VCM_M is approximately 2.3 times faster than both VCM_T and VCM_E . As explained earlier, VCM_M and VCM_T yield I/O costs similar to VCM_E . On average, VCM_M and VCM_T introduce 7.5% and 3.8% errors, respectively. These amounts of errors do not have noticeable impact on many practical applications. We observe from Eq. (1) that the cells generated near the target are very small. Therefore the approximations of the cells close to the target are almost similar to the corresponding exact cells. The cell size increases with the increase of distance and a significant portion of errors is introduced for these distant cells. As the visibility of the target from distant cells is insignificant in most of the cases, the error in such cells is tolerable.

In general, with the increase in ϑ , the introduced errors in VCM_M and VCM_T increase as a larger quad-block size approximates the cells with lesser accuracy. But with the increase in ϑ , the total processing time and I/O cost reduce significantly. Hence for applications that can tolerate reduced accuracy, a large ϑ can result into better performance.

5.4. Performance in 3D

In this section, we present the experimental results for 3D British and Boston datasets by varying different parameters.

5.4.1. Effect of μ

In this experiment, we vary the value of μ as 1, 2, 4, 8, and 16 arcmin and measure the total processing time and I/Os for British and Boston datasets (Fig. 10). We also present the errors resulted in the two approximation methods as shown in Table 6.

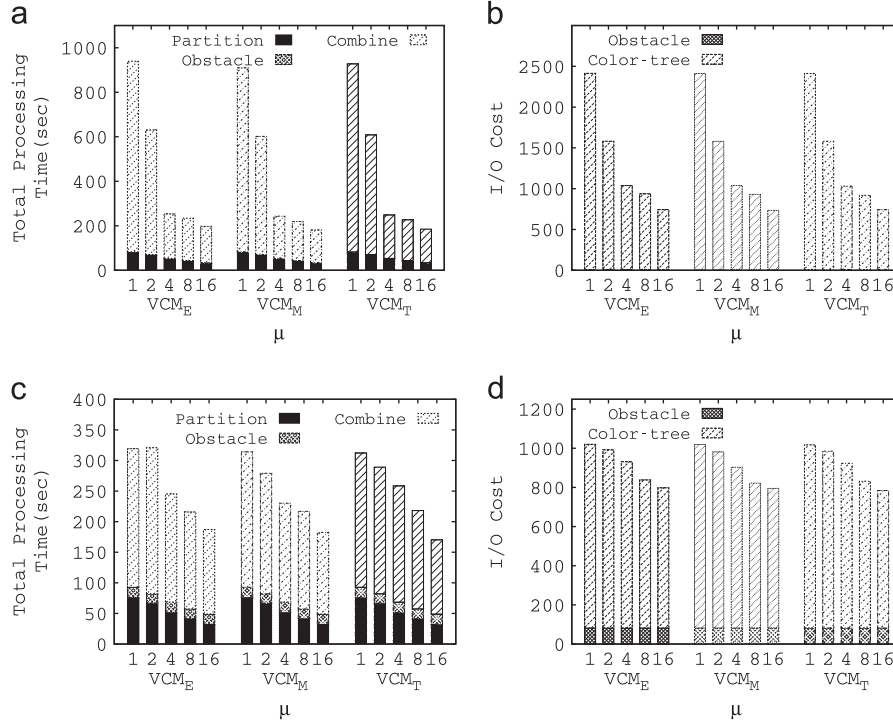


Fig. 10. Effect of μ in 3D British (a and b) and Boston (c and d).

Table 6

Errors in VCM_M and VCM_T for μ in 3D.

Dataset	Method	Error (%)				
		1	2	4	8	16
British	VCM_M	4.56	8.02	8.11	10.46	15.25
	VCM_T	4.1	7.63	7.76	8.89	13.89
Boston	VCM_M	3.31	5.03	8.02	9.12	10.05
	VCM_T	3.12	4.70	7.62	8.57	9.79

For British dataset, on average VCM_M and VCM_T run only 5% and 3% faster than VCM_E , respectively. The difference in I/O cost, i.e., the number of pages accessed is also negligible among the three methods. As discussed earlier in Section 3 the three methods differ only in the final phase, i.e., combining the outcome of previous two steps to produce the VCM. As we consider only full visibility of the target object from a partition cell, a huge number of cells in the dataspace fall into the obstructed region. So in the final phase the three solutions perform similarly as they have to combine the same color-tree and visible region quad-tree. The average errors introduced in VCM_M and VCM_T are 9% and 8%, respectively.

Although the results for both British and Boston datasets follow similar pattern, Boston dataset causes much smaller I/O and computational cost. The reason is, in the case of the densely populated Boston dataset, most of the obstacles are pruned rapidly during the visible region quad-tree formation. So there are lesser number of visibility computation and page access to construct VCM in

Boston dataset. The processing time of VCM_M is on average 8% faster than that of VCM_E and VCM_T . The average error introduced in VCM_M is 7%, whereas the more accurate VCM_T yields 6% error.

The I/O costs in the obstacle retrieval phase are the same in all three methods. Thus in subsequent sections, we show only the I/O cost required to combine the quad-tree and the color-tree (i.e., the final phase of constructing the VCM).

5.4.2. Effect of ϑ

In this experiment, we vary the quad-tree block size ϑ as 1, 2, 4, 8, and 16 times of the minimum threshold of a quad-tree block and show the results in Fig. 11 and Table 7.

For British dataset, the computational and I/O costs are similar for all three methods, e.g., on average VCM_M runs only 3% faster than VCM_E . Because, similar to the case of varying μ , the three phases of constructing VCM differ slightly for these methods. The average error introduced in VCM_M and VCM_T is about 10%. The results derived from Boston dataset are similar to that of British dataset. VCM_M runs about 3% faster than both VCM_E and VCM_T . The average errors introduced in the approximation methods are about 9%.

As the errors introduced by two approximations show similar trends for A_Q , FOV, and L_T , we show only the total processing time and the I/O cost in the subsequent sections.

5.4.3. Effect of A_Q

In this experiment, we vary the query area A_Q as 5%, 10%, 15%, 20%, and 25% of the total dataspace (Fig. 12(a)–(d)).

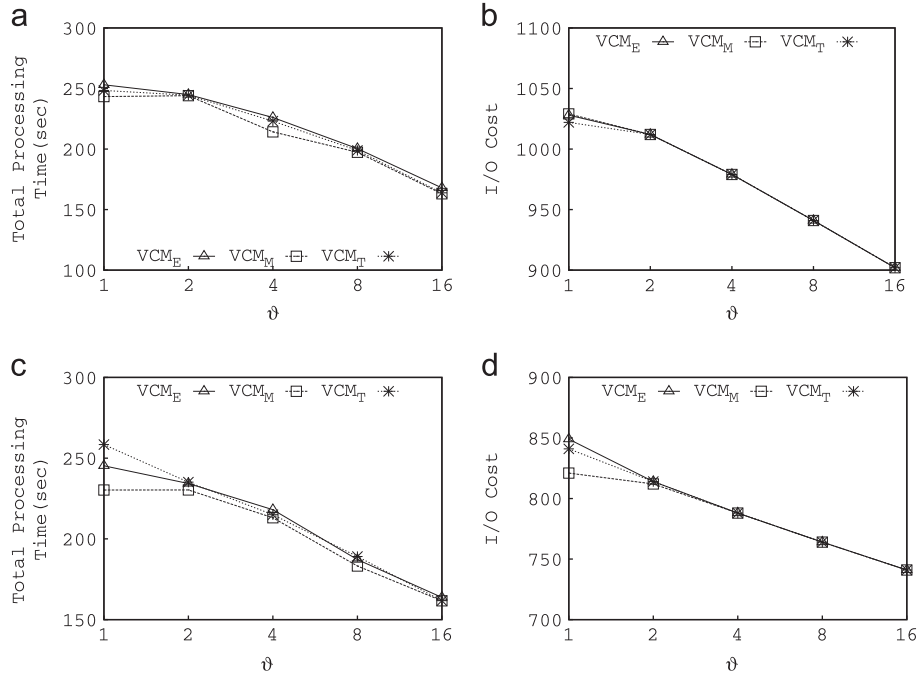


Fig. 11. Effect of θ in 3D British (a and b) and Boston (c and d).

Table 7
Errors in VCM_M and VCM_T for θ in 3D.

Dataset	Method	Error (%)				
		1	2	4	8	16
British	VCM_M	8.11	9.31	10.15	11.83	14.96
	VCM_T	7.76	9.00	9.86	11.05	12.72
Boston	VCM_M	8.02	9.08	9.38	9.82	11.03
	VCM_T	7.62	8.82	9.06	8.98	10.29

In the case of British dataset, VCM_M and VCM_T run 5% and 3% faster than VCM_E , respectively. For Boston dataset, VCM_M runs 5% faster than both VCM_E and VCM_T . The I/O cost and the total processing time increase for all three methods with the increase in A_Q , as we need to consider higher number of cells and obstacles to construct the VCM.

5.4.4. Effect of FOV

In this experiment, we vary the field of view (FOV) as 60, 120, 180, 240, 300, and 360 degrees for British and Boston datasets (Fig. 13). When FOV is set to 360°, the FOV covers the surrounding space in all directions. We observe that the processing time and I/Os increase with the increase of FOV, which is expected.

Like the previous cases, the total processing time and I/O cost are similar for all three methods. For British and Boston dataset, VCM_M runs 3% and 4% faster than both VCM_E and VCM_T , respectively. With the increase of FOV from 60 to 360°, the total processing time of VCM_E increases by nearly 40% and 46% for British and Boston dataset, respectively. On the other hand, the I/O costs

increase 33% and 31% for British and Boston dataset, respectively.

5.4.5. Effect of L_T

In this experiment, we vary the length of target L_T as 5%, 10%, 15%, 20%, and 25% of the total length of the dataspace. In general, with the change in L_T , no significant change in performance is observed for any of the datasets (Fig. 14).

5.4.6. Effect of varying D_S

In this experiment we vary the number and distribution of obstacles and measure the performance of our approximation methods in terms of I/O cost, total processing time, and approximation errors. We vary D_S as a set of 5k, 10k, 15k, 20k, and 25k obstacles while keeping the other parameters at their default values. We consider both Uniform (U) and Zipf (Z) distributions of the obstacles.

In general, as the number of obstacles increases, the area of the dataspace gets more obstructed. So during the final phase of VCM construction (i.e., combining color-tree and visible region quad-tree), computational costs and I/O costs get reduced. Consequently, the overall costs decrease with the increase in D_S for both Uniform and Zipf distribution. In the case of Uniform and Zipf distributions, as the number of obstacles is varied from 5k to 25k, the total processing time decreases nearly 40% and 20%, respectively for all three methods. On the other hand, the I/O costs decrease approximately 18% and 17% with the increase of D_S for Uniform and Zipf distributions, respectively.

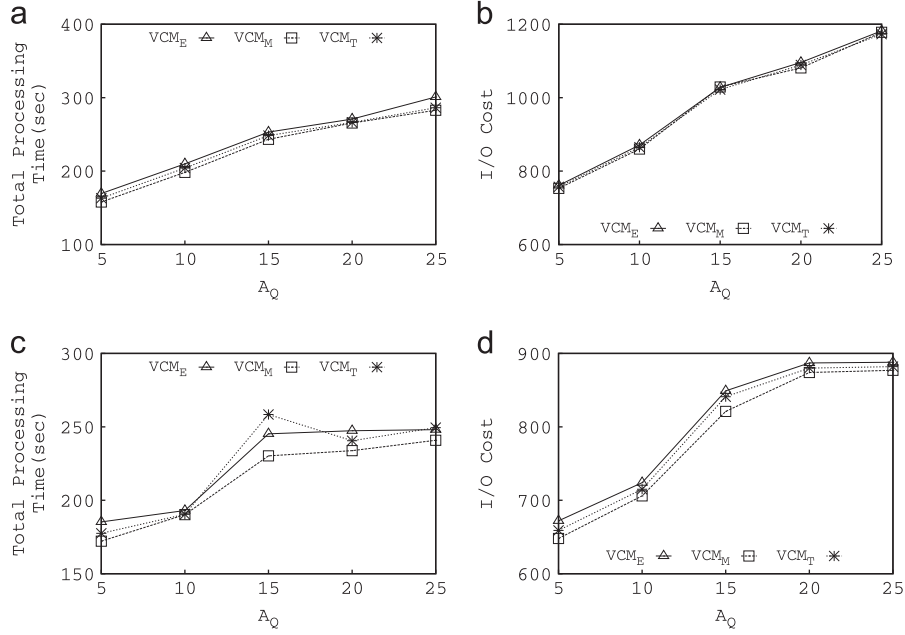


Fig. 12. Effect of A_Q in 3D British (a and b) and Boston (c and d).

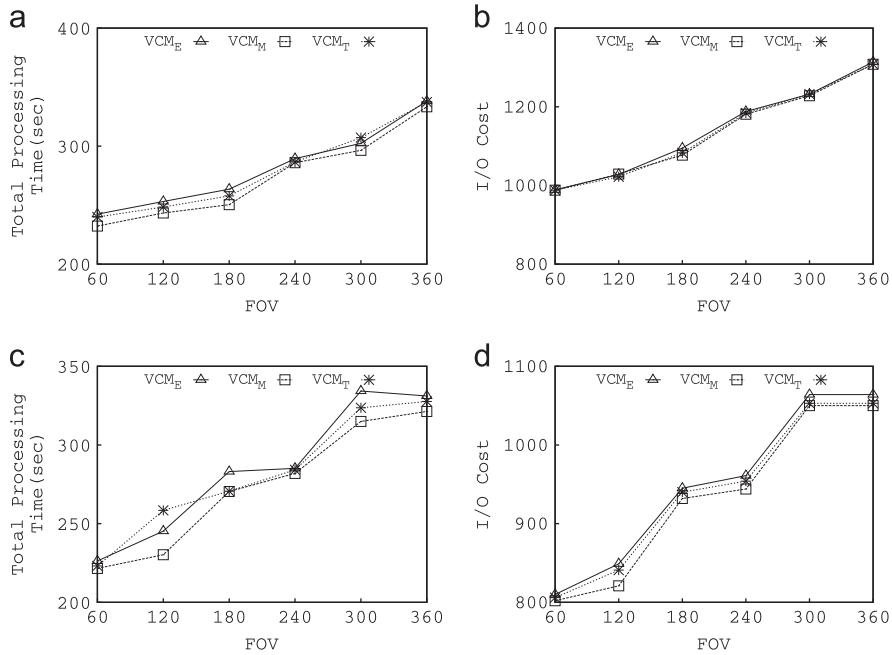


Fig. 13. Effect of FOV in 3D British (a and b) and Boston (c and d).

6. Related works

The notion of visibility is actively studied in different contexts: computational geometry (e.g., [21–23]), computer graphics and visualization (e.g., [4,9]), urban planning and architecture (e.g., [24–26]), and spatial databases (e.g., [1–3]). Most of these techniques consider visibility as a binary notion, i.e., a point is either visible or invisible from another point (Fig. 15).

6.1. Visibility in computational geometry

Several approaches have been proposed to construct visibility graph and visibility polygon [21–23,27–29] for visibility computation in the context of computational geometry. A visibility graph is defined by a set P of n points inside a polygon Q where two points $p, q \in P$ are joined by an edge if the segment $pq \subset Q$ [22].

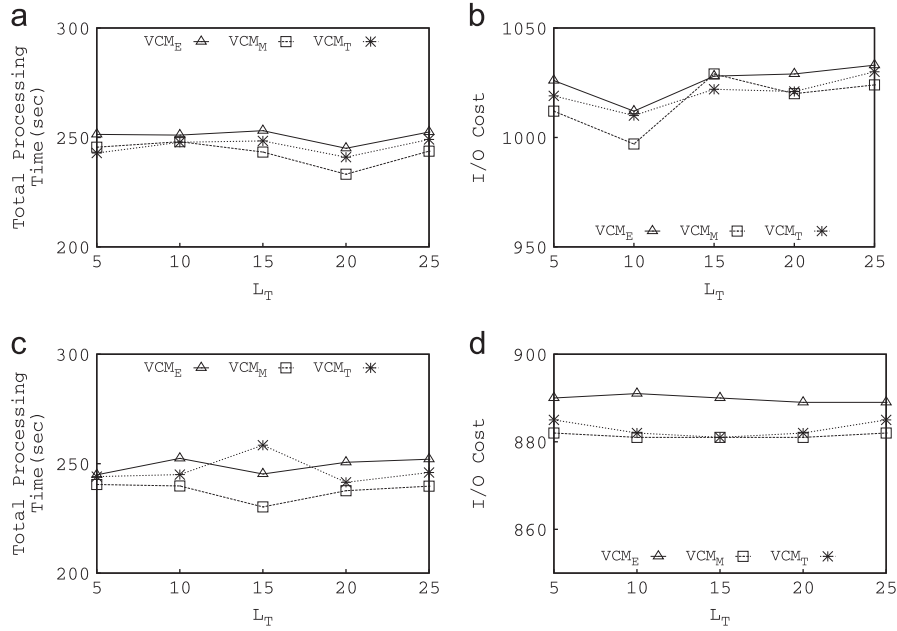


Fig. 14. Effect of L_T in 3D British (a and b) and Boston (c and d).

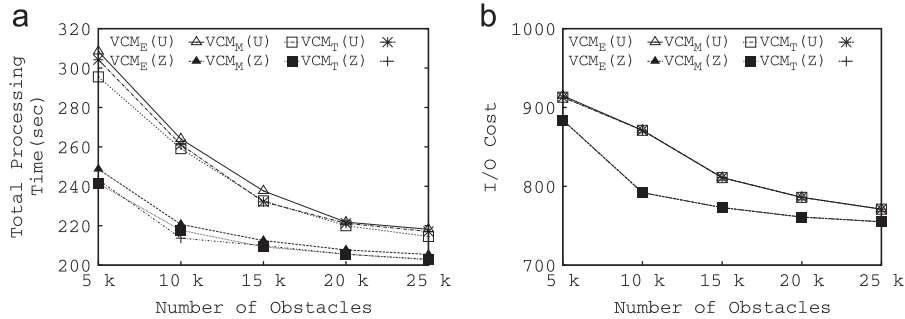


Fig. 15. Effect of varying dataset size in 3D (a and b).

The visibility polygon $V(q)$ of a point q in a polygon P is defined as the set of points in P that are visible from q [21]. Two points inside a polygon are visible to each other if their connecting segment remains completely inside the polygon. Zarei and Ghodsi consider the problem of computing the visibility polygon of a query point inside a non-simple polygon [21]. They require less time to compute the visibility polygon than the previous works, in exchange of costly preprocessing of the given input and an expensive data structure of size $O(n^3)$. Later, Asano et al. improved the preprocessing time and reduced the space requirement [23,29]. Similar studies are also conducted in [27,28].

Though the algorithms above efficiently solve the problem of constructing visibility polygon, they must rely on preprocessing and/or accessing all obstacles. Consequently, they are inappropriate for many spatial database applications as any update will invalidate the preprocessed data and accessing all objects for each query point is not feasible in such cases. Moreover, they consider visibility as a binary notion, which makes these approaches inapplicable where quantification of visibility is required.

6.2. Visibility in computer graphics and visualization

In computer graphics, the visibility map refers to a planar subdivision that encodes the visibility information, i.e., which points are mutually visible [4]. Two points are mutually visible if the straight line segment connecting these points does not intersect with any obstacle. If a scene is represented using a planar straight-line graph, a horizontal (vertical) visibility map is obtained by drawing a horizontal (vertical) straight line l through each vertex p of that graph until l intersects an edge e of the graph or extends to infinity. The edge e is horizontally (vertically) visible from p . A large body of works [5–9] construct such visibility maps efficiently.

Given a collection of surfaces representing boundaries of obstacles, Tsai et al. consider visibility problem as determining the regions of space or the surfaces of obstacles that are visible to an observer [30,31]. They model visibility as a binary notion and find the light and dark regions of a space for a point light source. Kim et al. introduced the completely visible region [32], which is

defined as the set of all points in 3D space that are completely visible from a given point S .

Above methods involve the computation of visible surfaces from a viewpoint and do not consider visibility factors such as angle and distance to quantify the visibility of a target object.

6.3. Visibility in urban planning and architecture

The visibility related problems are also studied in urban planning and architecture, often with the use of geographical information systems data as a modeling tool [33–35]. In environmental planning, the visibility is assessed using computerized simulation to predict the visual consequences of alterations in the urban landscapes [24,25]. Hernandez et al. [26] propose a method that is useful to planners and designers to choose locations where new buildings best integrate into the landscape in terms of its visual impact. These methods also do not quantify the visibility and treat visibility as a binary notion. Recently Bartie et al. [36] propose a number of visual metrics to quantify how much of a landmark is visible from its surroundings. This approach assumes all obstacles are in main memory. Moreover, the work in [36] does not prune irrelevant obstacles of the dataset while determining the visibility, and hence is not suitable for applications comprising a large number of obstacles.

Software systems for 3D data visualization and rendering are also available for architects and urban planner. Lumion [37] is an architectural visualization and walk-through software for urban planning, which focuses on scene creation and video rendering of 3D data. Google SketchUp [38] is widely used in 3D modeling for engineering applications and enables placement of its models in Google Earth [39]. Autodesk [40] provides a wide range of softwares including AutoCAD [41] for 3D modeling and Maya [42] for 3D animation and simulation. The other commercial softwares [43,44] used in urban planning and architecture mostly focus on 3D modeling and rendering. Though these softwares facilitate platforms to create and render 3D objects with functionalities like animation and walk-through, they do not provide any functionality for quantifying visibility of a 3D object. Our proposed VCM techniques can be integrated with these softwares to enable them answering many realistic visibility queries that require quantification of visibility in a 3D space.

6.4. Visibility in spatial queries

Visibility problems studied in spatial databases usually involve finding the nearest object to a given query point in the presence of obstacles. In recent years, several variants of the nearest neighbor (NN) queries have been proposed that include Visible Nearest Neighbor (VNN) query [1], Obstructed NN (ONN) query [45], Continuous Obstructed NN (CONN) query [2], and Continuous Visible NN query [3].

A CNN query finds the k NNs for a moving query point [46]. Gao et al. propose a variation of CNN; namely, a CONN query [2]. Given a dataset P , an obstacle set O , and a query line segment q in a 2D space, a CONN

query retrieves the NN of each point on q according to the obstructed distance. Nutanong et al. introduce an approach to find the NN that is visible to a query point [1], which uses an incremental approach to retrieve the obstacles efficiently. For efficient obstacle retrieval in 3D dataspace, many other algorithms are introduced recently [47–49].

The aforementioned spatial queries find the nearest object in an obstructed space from a given query point where query results are ranked according to visible distances from that query point. Instead of quantifying visibility as a non-increasing function from a target to a viewpoint, they label a particular point or region as either visible or invisible. But for constructing a VCM of the entire space, such binary notion is not applicable.

Recently the concept of maximum visibility query is tossed in [10] that considers the effect of obstacles during quantifying visibility of a target object. They measure the visibility of a target from a given set of query points and rank these query points based on the visibility measured as the visible surface area of the target. This approach is not applicable for constructing the VCM as we need to compute the visibility for every viewpoint of the dataspace by considering both the distance and the angle between a target and a viewpoint.

7. Conclusion

In this paper, we have proposed a technique to compute a visibility color map (VCM) that forms the basis of many real-life visibility queries in 2D and 3D spaces. A VCM quantifies the visibility of (from) a target object from (of) each viewpoint of the surrounding space and assigns colors accordingly in the presence of obstacles. Our approach exploits the limitation of a human eye or a lens to partition the space into cells in the presence of obstacles such that the target appears same from all viewpoints inside a cell.

Our proposed solution significantly outperforms the baseline approach by at least 800 times and six orders of magnitude in terms of computational time and I/O cost, respectively for both datasets. Our conducted experiments on real 2D and 3D datasets demonstrate the efficiency of our approximations, VCM_M and VCM_T . On average, the approximations VCM_M and VCM_T improve the processing time by 65% and 17% over VCM_E , respectively and require almost similar I/O costs. Both VCM_M and VCM_T improve the efficiency of VCM_E by introducing only 9% and 5% error on average.

Acknowledgments

This research has been conducted at Samsung Innovation Lab in the Department of Computer Science and Engineering (CSE), Bangladesh University of Engineering and Technology (BUET). This work is supported by the grant for Advanced Research in Science from the Ministry of Education, Govt. of the People's Republic of Bangladesh.

References

- [1] S. Nutanong, E. Tanin, R. Zhang, Incremental evaluation of visible nearest neighbor queries, *IEEE Trans. Knowl. Data Eng.* 22 (2010) 665–681.
- [2] Y. Gao, B. Zheng, Continuous obstructed nearest neighbor queries in spatial databases, in: *SIGMOD*, 2009, pp. 577–590.
- [3] Y. Gao, B. Zheng, W.-C. Lee, G. Chen, Continuous visible nearest neighbor queries, in: *EDBT*, 2009, pp. 144–155.
- [4] Algorithms and Theory of Computation Handbook, CRC Press LLC, 1999.
- [5] J. Bittner, Efficient construction of visibility maps using approximate occlusion sweep, in: *SCCG*, 2002, pp. 167–175.
- [6] J. Grasset, O. Terraz, J. Hasenfratz, D. Plemenos, et al., Accurate scene display by using visibility maps, in: *SCCG*, 1999.
- [7] T. Keeler, J. Fedorkiw, S. Ghali, The spherical visibility map, *Comput. Aided Des.* 39 (1) (2007) 17–26.
- [8] P. Wonka, Visibility in computer graphics, *EPB Plan. Des.* 30 (2003) 729–755.
- [9] A. Stewart, T. Karkanis, Computing the approximate visibility map, with applications to form factors and discontinuity meshing, in: *EGWR*, 1998, pp. 57–68.
- [10] S. Masud, F.M. Choudhury, M.E. Ali, S. Nutanong, Maximum visibility queries in spatial databases, in: *ICDE*, 2013, pp. 637–648.
- [11] C. Rao, E. Wegman, J. Solka, *Handbook of Statistics*, vol. 24, North-Holland Publishing Co, 2005.
- [12] J. Baird, *Psychophysical Analysis of Visual Space*, Pergamon Press, London, 1970.
- [13] D. McCready, On size, distance, and visual angle perception, *Atten. Percept. Psychophys.* 37 (4) (1985) 323–334.
- [14] P. Kaiser, *The Joy of Visual Perception*, York University, 1996.
- [15] K. Morling, *Geometric and Engineering Drawing*, 3rd edition, Routledge, 2012.
- [16] B.H. Walker, *Optical Engineering Fundamentals*, 2nd ed. SPIE Press, 2009.
- [17] N. Beckmann, H.-P. Kriegel, R. Schneider, B. Seeger, The R*-tree: an efficient and robust access method for points and rectangles, in: *SIGMOD*, 1990, pp. 322–331.
- [18] A. Guttman, R-trees: a dynamic index structure for spatial searching, in: *SIGMOD*, 1984, pp. 47–57.
- [19] D. Atchison, G. Smith, *Optics of the Human Eye*, Butterworth-Heinemann, Oxford, UK, 2000.
- [20] H. Samet, *The Design and Analysis of Spatial Data Structures*, Addison-Wesley, MA, 1990.
- [21] A.R. Zarei, M. Ghodsi, Efficient computation of query point visibility in polygons with holes, in: *SCG*, 2005, pp. 6–8.
- [22] B. Ben-Moshe, O. Hall-Holt, M.J. Katz, J.S.B. Mitchell, Computing the visibility graph of points within a polygon, in: *SCG*, 2004, pp. 27–35.
- [23] T. Asano, T. Asano, L.J. Guibas, J. Hershberger, H. Imai, Visibility of disjoint polygons, in: *Algorithmica*, 1986, pp. 49–63.
- [24] E. Lange, Integration of computerized visual simulation and visual assessment in environmental planning, *Landsc. Urb. Plan.* 30 (1–2) (1994) 99–112.
- [25] H.-D. Shang, A method for creating precise low-cost landscape architecture simulations combining computer-aided design with computer video-imaging techniques, *Landsc. Urb. Plan.* 22 (1992) 11–16.
- [26] J. Hernandez, L. Garca, F. Ayuga, Assessment of the visual impact made on the landscape by new buildings: a methodology for site selection, vol. 68, 2004, pp. 15–28.
- [27] S. Suri, J. O'Rourke, Worst-case optimal algorithms for constructing visibility polygons with holes, in: *Symposium on Computational Geometry*, 1986, pp. 14–23.
- [28] P.J. Heffernan, J.S.B. Mitchell, An optimal algorithm for computing visibility in the plane, in: *WADS*, 1991, pp. 437–448.
- [29] T. Asano, T. Asano, L.J. Guibas, J. Hershberger, H. Imai, Visibility-polygon search and Euclidean shortest paths, in: *FOCS*, 1985, pp. 155–164.
- [30] Y. Tsai, L. Cheng, S. Osher, P. Burchard, G. Sapiro, Visibility and its dynamics in a PDE based implicit framework, *J. Comput. Phys.* 199 (1) (2004) 260–290.
- [31] L. Cheng, Y. Tsai, Visibility optimization using variational approaches, *Commun. Math. Sci.* 3 (3) (2005) 425–451.
- [32] D.S. Kim, K.-H. Yoo, K.-Y. Chwa, S.Y. Shin, Efficient algorithms for computing a complete visibility region in three-dimensional space, in: *Algorithmica*, 1998, pp. 201–225.
- [33] E. Peccol, C.A. Bird, T.R. Brewer, GIS as a tool for assessing the influence of countryside designations and planning policies on landscape change, *J. Environ. Manag.* 47 (1996) 355–367.
- [34] W. Suleiman, T. Joliveau, E. Favier, 3D urban visibility analysis with vector GIS data, in: *GISRUK*, 2011.
- [35] P.P.-J. Yang, S.Y. Putra, W. Li, Viewsphere: a gis-based 3D visibility analysis for urban design evaluation, in: *Environment and Planning B: Planning and Design*, vol. 34, 2007.
- [36] P. Bartie, F. Reitsma, S. Kingham, S. Mills, *Advancing Visibility Modelling Algorithms for Urban Environments*, vol. 34, 2010, pp. 518–531.
- [37] Lumion, (<http://lumion3d.com/urban-planning-software>).
- [38] Google Sketchup, (<http://www.sketchup.com>).
- [39] Google Earth, (<http://www.google.com/earth>).
- [40] Autodesk, (<http://www.autodesk.com>).
- [41] Autocad, (<http://www.autodesk.com.au/products/autodesk-autocad/overview>).
- [42] Maya, (<http://www.autodesk.com/products/autodesk-maya/overview>).
- [43] Cityengine, (<http://www.esri.com/software/cityengine>).
- [44] Virtuelcity, (<http://www.virtuelcity.com>).
- [45] J. Zhang, D. Papadias, K. Mouratidis, M. Zhu, Spatial queries in the presence of obstacles, in: *EDTB*, 2004, pp. 366–384.
- [46] Y. Tao, D. Papadias, Q. Shen, Continuous nearest neighbor search, in: *Vldb*, 2002, pp. 287–298.
- [47] I. Sipiran, B. Bustos, T. Schreck, Data-aware 3D partitioning for generic shape retrieval, in: *Computers and Graphics*, 2013, pp. 460–472.
- [48] P. Papadakis, I. Pratikakis, T. Theoharis, G. Passalis, S. Perantonis, 3D object retrieval using an efficient and compact hybrid shape descriptor, in: *Eurographics Workshop on 3D Object Retrieval*, 2008.
- [49] M. Ali, E. Tanin, R. Zhang, L. Kulik, A motion-aware approach for efficient evaluation of continuous queries on 3D object databases, *Vldb J.* 19 (5) (2010) 603–632.