# Robot Navigation for Automatic Model Construction using Safe Regions 

Héctor González-Baños Jean-Claude Latombe<br>Department of Computer Science<br>Stanford University, Stanford, CA 94305, USA<br>e-mail:\{hhg,latombe\}@robotics.stanford.edu


#### Abstract

Automatic model construction is a core problem in mobile robotics. To solve this task efficiently, we need a motion strategy to guide a robot equipped with a range sensor through a sequence of "good" observations. Such a strategy is generated by an algorithm that repeatedly computes locations where the robot must perform the next sensing operation. This is called the next-best view problem. In practice, however, several other considerations must be taken into account. Of these, two stand out as decisive. One is the problem of safe navigation given incomplete knowledge about the robot surroundings. The second one is the issue of guaranteeing the alignment of multiple views, closely related to the problem of robot selflocalization. The concept of safe region proposed in this paper makes it possible to simultaneously address both problems.


## 1. Introduction

Automatic model construction is a fundamental task in mobile robotics [1]. The basic problem is easy to formulate: After being introduced into an unknown environment, a robot, or a team of robots, must perform sensing operations at multiple locations and integrate the acquired data into a representation of the environment. Despite this simple formulation, the problem is difficult to solve in practice. First, there is the problem of choosing an adequate representation of the environment - e.g., topological maps [2], polygonal layouts [1], occupancy grids [4], 3-D models [12], or featurebased maps [6]. Second, the representation must be extracted from imperfect sensor readings - e.g., depth readings from range-sensors may fluctuate due to changes in surface textures [3], different sets of 3-D scans must be zippered [13], and captured images must be aligned and registered [11]. Finally, if the system is truly automatic, the robot must decide on its own the necessary motions to construct the model [5].

Past research in model construction has mainly focused on developing techniques for extracting relevant features (e.g., edges, corners) from raw sensor data, and on integrating these into a single and consistent model. There is also prior research on the computation of sensor motions, mostly on finding the next-best view (NBV) [3, 11]: Where should the sensor be placed for the next sensing operation? Typically, a model is first built by combining images taken from a few distributed viewpoints. The resulting model usually contains gaps. An NBV technique is then used to select additional viewpoints that will provide the data needed to fill the remaining gaps.

Traditional NBV approaches are not suitable for mobile robotics. One reason is that most of the existing NBV techniques have been designed for systems that build a 3-D model of a relatively small object using a precise range sensor moving around the specimen. Collisions, however, are not a major issue for sensors that are mechanically constrained to operate outside the convex hull of the scene. In robotic applications, by contrast, the sensor navigates within the convex hull of the scene. Therefore, safe navigation considerations must always be taken into account when computing the next-best view for a robot map builder.

The second reason why most existing NBV techniques cannot be applied to mobile robots is that very few of the proposed approaches explicitly consider imageregistration issues (one exception is the sensor-based technique presented in [11]). Localization problems particularly affect mobile sensors, and image registration becomes paramount when it is the means by which a mobile robot re-localizes itself (this is the so-called simultaneous localization and map building problem) [9, 7]. Although many image-registration techniques can be found in the literature, all require that each new image significantly overlaps with portions of the environment seen by the robot at previous sensing locations [9].

The system presented in [5] deals with the safe navigation and localization problems by applying the concept of safe region and the NBV algorithm introduced in this paper. With safe regions, it is possible to iteratively build a map by executing union operations over successive views, and use this map for motion planning. Moreover, safe regions can be used to estimate the overlap between future views and the current global map, and to compute locations that could potentially see unexplored areas.

The work in [5] is mainly about system integration and proof of concept. Instead, this paper focuses on the formal definition of a safe region (Section 2), and describes how to compute such region from sensor data (Section 3). An NBV algorithm based on safe regions is outlined in Section 5, and Section 6 describes an experimental run using our system.

## 2. Definition of Safe Regions

Suppose that the robot is equipped with a polar range sensor measuring the distance from the sensor's center to objects lying in a horizontal plane located at height $h$ above the floor. Because all visual sensors are limited in range, we assume that objects can only be detected within a distance $d_{M}$. In addition, most range-finders cannot reliably detect surfaces oriented at grazing angles with respect to the sensor. Hence, we also assume that surface points that do not satisfy the sensor's incidence constraint cannot be reliably detected by the sensor. Formally, our visibility model is the following:
Definition 2.1 (Visibility under Incidence and Range Constraints) Let the open subset $\mathcal{W} \subset \Re^{2}$ describe the workspace layout. Let $\partial \mathcal{W}$ be the boundary $\mathcal{W}$. A point $w \in \partial \mathcal{W}$ is said to be visible from $q \in \mathcal{W}$ if the following conditions are true:

1. Line of sight constraint: The segment from $q$ to $w$ doesn't intersect $\partial \mathcal{W}$.
2. Range constraint: $d(q, w) \leq d_{M}$, where $d(q, w)$ is the Euclidean distance between $q$ and $w$, and $d_{M}>0$ is an input constant.
3. Incidence constraint: $\angle(\boldsymbol{n}, \boldsymbol{v}) \leq \tau$, where $\boldsymbol{n}$ is a vector perpendicular to $\partial \mathcal{W}$ at $w, \boldsymbol{v}$ is oriented from $w$ to $q$, and $\tau \in[0, \pi / 2]$ is an input constant.


Figure 1. Effect of incidence on safe regions.
Without any loss of generality we assume the sensor is located at the origin (the workspace can always be re-mapped to a reference frame centered on the sensor). The sensor's output is assumed to be as follows:
Definition 2.2 (Range Sensor Output) The output of a range sensor is an ordered list $\Pi$, representing the sections of $\partial \mathcal{W}$ visible from the origin under Definition 2.1. Every $r(\theta ; a, b) \in \Pi$ is a polar function describing a section of $\partial \mathcal{W}$, and such function is continuous $\forall \theta \in(a, b)$ and undefined elsewhere. $\Pi$ contains at most one function defined for any $\theta \in(-\pi, \pi]$ (i.e., no two functions overlap), and the list is ordered counter-clockwise.

Given an observation $\Pi$ made by the robot at a location $q$, we define the local safe region $s_{l}$ at $q$ as the largest region guaranteed to be free of obstacles. While range restrictions have an obvious impact on $s_{l}$, the effect of incidence is more subtle. In Figure $1(a)$, a sensor detects the surface contour shown in black. A naive approach may construct the region in light color (yellow) by joining the detected surfaces with the perimeter limit of the sensor, and consider this region free from obstacles. Because the sensor is unable to detect surfaces oriented at grazing angles, this region may be not be safe, as shown in (b). A true safe region is shown in (c), for an incidence constraint of $\tau=70 \mathrm{deg}$.

## 3. Computing Safe Regions

The region $s_{l}$ is bounded by solid and free curves. A solid curve represents an observed section of $\partial \mathcal{W}$, and is contained in the list $\Pi$. Given two solid curves $\left\{r_{1}\left(\theta ; a_{1}, b_{1}\right), r_{2}\left(\theta ; a_{2}, b_{2}\right)\right\} \subseteq \Pi, r_{2}$ is said to succeed $r_{1}$ if no other element in $\Pi$ is defined in the interval $\left[b_{1}, a_{2}\right]$. A curve $f\left(\theta ; b_{1}, a_{2}\right)$ joining a pair $\left(r_{1}, r_{2}\right)$ of successive sections is called a free curve if: (1) no undetected obstacle is contained in the polar region $b_{1}<\theta<a_{2}$ bounded by $f$; and (2) this region is the largest possible.

The main goal of this section is to find the free curves that join each successive pair in $\Pi$ in order to bound the region $s_{l}$. It turns out that the complexity of $f$ is $\mathrm{O}(1)$. In fact, a free curve $f$ can be described using no more than 3 function primitives:
Theorem 3.1 (Free Curves) Suppose $r_{2}\left(\theta ; a_{2}, b_{2}\right)$ succeeds $r_{1}\left(\theta ; a_{1}, b_{1}\right)$ in the output list $\Pi$. If $\partial \mathcal{W}$ is continuously differentiable, then the free curve $f\left(\theta ; b_{1}, a_{2}\right)$ connecting $r_{1}$ to $r_{2}$ consists of at most three pieces. Each piece is either a line segment, a circular arc, or a section of a logarithmic spiral.

The rest of this section proves this claim. But first we need the following lemma:

Lemma 3.2 Let $r_{2}\left(\theta ; a_{1}, b_{1}\right)$ succeed $r_{2}\left(\theta ; a_{2}, b_{2}\right)$ in the list $\Pi$. Let $C$ be some obstacle, and suppose that neither $r_{1}$ nor $r_{2}$ are part of the boundary of $C$ (i.e., $C$ is disjoint from $r_{1}$ and $r_{2}$ ). If $\partial \mathcal{W}$ is continuously differentiable, then no portion of $C$ lies within a distance $d_{M}$ from the origin in the polar interval $b_{1}<\theta<a_{2}$.
Proof: Suppose the lemma is not true - that is, there is a portion of $C$ within $d_{M}$ of the origin inside the polar interval $\left(b_{1}, a_{2}\right)$. Let $p$ be the closest point to the origin in the boundary of $C$. Because $\partial \mathcal{W}$ is differentiable, the normal of $\partial \mathcal{W}$ at $p$ points toward the origin. Therefore, $p$ and its vicinity should have been observed. The vicinity of $p$ must then be part of an element of $\Pi$. But this contradicts our assumption that $r_{2}$ succeeds $r_{1}$ and that $C$ is disjoint from $r_{1}$ and $r_{2}$.

From here on, let $\beta=a_{2}-b_{1}, \rho_{1}=r_{1}\left(b_{1}\right)$ and $\rho_{2}=r_{2}\left(a_{2}\right)$; and let $l_{1}$ and $l_{2}$ denote the rays joining the origin with $p_{1}=\left(\rho_{1}, b_{1}\right)$ and $p_{2}=\left(\rho_{2}, a_{2}\right)$, respectively.

Each endpoint of a curve in $\Pi$ represents one of the following events: the sensor line-of-sight was occluded (denoted as case $\{0\}$ ), the range constraint was exceeded (case $\{\mathrm{e}\}$ ), or the incidence constraint was exceeded (case $\{\mathrm{v}\}$ ). To join $p_{1}$ with $p_{2}$ there are a total of 6 distinct cases: $\{\mathrm{v}, \mathrm{v}\},\{\mathrm{v}, \mathrm{o}\},\{\mathrm{v}, \mathrm{e}\},\{\mathrm{e}, \mathrm{e}\},\{\mathrm{o}, \mathrm{o}\}$ and $\{\mathrm{e}, \mathrm{o}\}$. The cases $\{\mathrm{o}, \mathrm{e}\},\{\mathrm{o}, \mathrm{v}\}$ and $\{\mathrm{e}, \mathrm{v}\}$ are mirror images of other cases.
Case $\{\mathbf{v}, \mathbf{v}\}$ : The incidence constraint was exceeded at $\theta=b_{1}$ and $\theta=a_{2}$. Therefore, the normal to $\partial \mathcal{W}$ immediately after $r_{1}$, and immediately before $r_{2}$, is oriented at a grazing angle with respect to the sensor. Suppose that $\partial \mathcal{W}$ continues after $r_{1}$ with its surface normal constantly oriented at exactly an angle $\tau$ with respect to the sensor's line-of-sight. This curve in polar coordinates satisfies the following relations:

$$
\begin{align*}
\boldsymbol{n} & \doteq-r \delta \theta \hat{\boldsymbol{e}}_{r}+\delta r \hat{\boldsymbol{e}}_{\theta}  \tag{1}\\
\boldsymbol{n} \cdot\left(-r \hat{\boldsymbol{e}}_{r}\right)=r|\boldsymbol{n}| \cos (\tau) & \Longrightarrow \frac{1}{r} \frac{\delta r}{\delta \theta}= \pm \lambda, \text { with } \lambda \doteq \tan (\tau) \tag{2}
\end{align*}
$$

Hence, the curve's equation is $r=r_{o} \exp \left[ \pm \lambda\left(\theta-\theta_{o}\right)\right]$, with $r_{o}=\rho_{1}$ and $\theta_{o}=b_{1}$. The equation now defines two spirals: a spiral $s_{1}^{+}$growing counter-clockwise from $p_{1}$ (or shrinking clockwise), and a second spiral $s_{1}^{-}$shrinking counter-clockwise from $p_{1}$ (or growing clockwise). $\partial \mathcal{W}$ must continue counter-clockwise from $p_{1}$ either "above" $s_{1}^{+}$or "below" $s_{1}^{-}$; otherwise, the incidence constraint would not have been violated.

Similarly, for the opposite end $p_{2}$, let $r_{o}=\rho_{2}$ and $\theta_{o}=a_{2}$. The solution to equation (2) now defines a spiral $s_{2}^{-}$growing clockwise from $p_{2}$ (or shrinking counterclockwise), and a second spiral $s_{2}^{+}$shrinking clockwise from $p_{2}$ (or growing counterclockwise). $\partial \mathcal{W}$ must continue clockwise from $p_{2}$ either "above" $s_{2}^{-}$or "below" $s_{2}^{+}$.

Remark 1. $\partial \mathcal{W}$ cannot continue below $s_{1}^{-}$when $\rho_{1} \exp (-\lambda \beta)<\rho_{2}$. In other words, $\partial \mathcal{W}$ cannot continue below $s_{1}^{-}$if this spiral curve cuts $l_{2}$ below the point $p_{2}$ (Figure $2(a)$ ). To show this, suppose $\partial \mathcal{W}$ continues below $s_{1}^{-}$, which implies that $\partial \mathcal{W}$ bends toward the sensor immediately after $r_{1}$. We know that $\partial \mathcal{W}$ does not cross the origin, else nothing is visible under Definition 2.1 and $\Pi$ would be empty. Hence, $\partial \mathcal{W}$ would have to bend outwards before cutting the ray $l_{2}$, otherwise $r_{2}$ will be occluded. Since $\partial \mathcal{W}$ is differentiable, there must then be a point $p$ where the normal to $\partial \mathcal{W}$ points towards the origin. Because of Lemma 3.2, this point $p$ is not occluded by any other section of $\partial \mathcal{W}$ that is disjointed from $r_{1}$ and $r_{2}$. Therefore, the vicinity of $p$ is a visible portion of $\partial \mathcal{W}$. This violates our assumption that $r_{2}$ succeeds $r_{1}$. Thus, when $\rho_{1} \exp (-\lambda \beta)<\rho_{2}$, the first section of the curve $f$ joining $r_{1}$ to $r_{2}$ coincides with $s_{1}^{+}$.

(a)

(b)

(c)

Figure 2. Example of a free-curve construction: (a) this situation is impossible; (b) in this case the free curve is composed of the segment joining $p_{1}$ with $p$ and the spiral $s_{2}^{-}$joining $p$ with $p_{2} ;(c)$ here the free curve is composed of the spiral $s_{1}^{+}$joining $p_{1}$ with $p$ and the spiral $s_{2}^{-}$joining $p$ with $p_{2}$ (unless $p$ is beyond range, in which case a circular arc of radius $d_{M}$ is added).

Remark 2. By symmetry, when $\rho_{2} \exp (-\lambda \beta)<\rho_{1}$ (i.e., $s_{2}^{+}$cuts $l_{1}$ below $p_{1}$ ), the last section of the curve $f$ coincides with $s_{2}^{-}$(which grows clockwise from $p_{2}$ ).

The point $p_{2}$ may lie below the intersection of $s_{1}^{-}$with $l_{2}$, above the intersection of $s_{1}^{+}$with $l_{2}$, or between both intersections. Likewise, the point $p_{1}$ may lie below the intersection of $s_{2}^{+}$with $l_{1}$, above the intersection of $s_{2}^{-}$with $l_{1}$, or between both intersections. There are total of 9 combinations of events for case $\{v, v\}$, but only 3 of them are independent:
(a) $s_{1}^{-}$cuts $l_{2}$ above $p_{2}$. Thus, $\rho_{1} \exp (-\lambda \beta)>\rho_{2}$, and this is equivalent to $\rho_{2} \exp (\lambda \beta)<\rho_{1}$. That is, $s_{2}^{-}$cuts $l_{1}$ below $p_{1}$.
(b) $s_{1}^{+}$cuts $l_{2}$ below $p_{2}$. Thus, $\rho_{1} \exp (\lambda \beta)<\rho_{2}$, and this is equivalent to $\rho_{2} \exp (-\lambda \beta)>\rho_{1}$. That is, $s_{2}^{+}$cuts $l_{1}$ above $p_{1}$.
(c) $s_{1}^{-}$cuts $l_{2}$ below $p_{2}$ and $s_{1}^{+}$cuts $l_{2}$ above $p_{2}$. Thus, $\rho_{1} \exp (-\lambda \beta)<$ $\rho_{2}<\rho_{1} \exp (\lambda \beta)$, and this is equivalent to $\rho_{2} \exp (-\lambda \beta)<\rho_{1}<$ $\rho_{2} \exp (\lambda \beta)$. That is, $s_{2}^{+}$cuts $l_{1}$ below $p_{1}$ and $s_{2}^{-}$cuts $l_{1}$ above $p_{1}$.
Let us analyze the first situation. $\rho_{1} \exp (-\lambda \beta)>\rho_{2}$ is equivalent to $\rho_{2} \exp (\lambda \beta)<\rho_{1}$, which in turn implies that $\rho_{2} \exp (-\lambda \beta)<\rho_{1}$. In other words, both the clockwise-growing $s_{2}^{-}$and the clockwise-shrinking $s_{2}^{+}$cut $l_{1}$ below $p_{1}$ (see Figure 2(b)). From Remark 2, the last section of the free curve $f$ coincides with $s_{2}^{-}$. Let $p$ be the intersection between $s_{2}^{-}$and $l_{1}$. The free curve $f$ joining $r_{1}$ to $r_{2}$ is thus composed of the segment joining $p_{1}$ with $p$ and the spiral $s_{2}^{-}$joining $p$ with $p_{2}$.

A symmetric argument applies to the second situation, when $\rho_{2} \exp (-\lambda \beta)>\rho_{1}$ (i.e., $s_{2}^{+}$cuts $l_{1}$ above $p_{1}$ ), except that Remark 1 is used in this case.

The only remaining situation is $(c): \rho_{1} \exp (-\lambda \beta)<\rho_{2}$ and $\rho_{2} \exp (-\lambda \beta)<\rho_{1}$. From Remarks 1 and 2, these inequalities imply that the first section of $f$ coincides with $s_{1}^{+}$while the last section of $f$ coincides with $s_{2}^{-}$. Let $p$ be the intersection of $s_{1}^{+}$ and $s_{2}^{-}$. If $p$ is within $d_{M}$, then $f$ is composed of the spiral $s_{1}^{+}$joining $p_{1}$ with $p$ and the spiral $s_{2}^{-}$joining $p$ with $p_{2}$ (Figure $2(c)$ ). Otherwise $p$ is beyond range, and $f$ is composed of a section of $s_{1}^{+}$, a circular arc of radius $d_{M}$, and a section of $s_{2}^{-}$.


Figure 3. Dealing with corners: $(a)$ the normal to $\partial \mathcal{W}$ at a corner is generalized as the average of $\boldsymbol{n}^{+}$and $\boldsymbol{n}^{-} ;(b)$ if we assume that a corner has "thickness", and is therefore detectable by the sensor, then any wedge-shaped object is visible if $\tau \geq 45 \mathrm{deg}$.

Case $\{\mathbf{v}, \mathbf{o}\}$ : As in the previous case, the curve $r_{1}$ was interrupted at $\theta=b_{1}$ because the incidence constraint was exceeded. The curve $r_{2}$, however, was interrupted at $\theta=a_{2}$ because a portion of $\partial \mathcal{W}$ blocked the sensor's line-of-sight. In order to produce the occlusion, $\partial \mathcal{W}$ must be tangent to $l_{2}$ at some point $p_{t}$ below $p_{2}$. We know from Lemma 3.2 that the portion of $\partial \mathcal{W}$ producing the occlusion cannot be disjointed from $r_{1}$. Thus, $p_{t}$ is part of the same curve as $r_{1}$.
$\partial \mathcal{W}$ cannot continue from $r_{1}$ below $s_{1}^{-}$. To show this, suppose $\partial \mathcal{W}$ continues below $s_{1}^{-}$. This implies that $\partial \mathcal{W}$ bends toward the sensor immediately after $r_{1}$. But to cause the occlusion, $\partial \mathcal{W}$ has to bend outwards before it reaches the tangent point $p_{t}$. Since $\partial \mathcal{W}$ is differentiable, there must be a point where the normal to $\partial \mathcal{W}$ points towards the origin. But we already know that this violates our assumption that $r_{2}$ succeeds $r_{1}$. Therefore, $\partial \mathcal{W}$ must continue above $s_{1}^{+}$.

For case $\{\mathrm{v}, \mathrm{o}\}$, it is always true that $s_{1}^{+}$cuts the ray $l_{2}$ below $p_{2}$ at some point $p$. Otherwise, it will be impossible to produce the occlusion at $p_{t}$, because $\partial \mathcal{W}$ continues from $r_{1}$ above $s_{1}^{+}$. Thus, $f$ is composed of the spiral $s_{1}^{+}$joining $p_{1}$ with $p$, and the segment joining $p$ with $p_{2}$.
Case $\{\mathbf{v}, \mathbf{e}\}$ : As before, the incidence constraint was exceeded at $\theta=b_{1}$, but $r_{2}$ was interrupted because the range constraint was exceeded at $\theta=a_{2}$. That is, $\rho_{2}=d_{M}$.

The point $p_{1}$ is within range, hence $\rho_{1} \exp (-\lambda \beta)<\rho_{2}$ because $\rho_{2}=d_{M}$. This is exactly the situation described in Remark 1 of case $\{\mathrm{v}, \mathrm{v}\}$. Thus, $\partial \mathcal{W}$ cannot continue below $s_{1}^{-}$, and the first section of $f$ coincides with $s_{1}^{+}$.

If $s_{1}^{+}$cuts the ray $l_{2}$ below $p_{2}$ at some point $p$, then $f$ is composed of the spiral $s_{1}^{+}$joining $p_{1}$ with $p$, and the segment joining $p$ with $p_{2}$. Otherwise $p$ is beyond range, and $f$ is composed of a section of $s_{1}^{+}$and a circular arc of radius $d_{M}$.
Case $\{\mathbf{e}, \mathbf{e}\}: \quad$ This case is trivial. The free curve is a circular arc joining $p_{1}$ to $p_{2}$.
Cases $\{\mathbf{0}, \mathbf{o}\}$ and $\{\mathbf{e}, \mathbf{o}\}$ : The reader may verify that these cases are impossible by following the same line of reasoning used throughout this proof. We skip the details for lack of space.

We have accounted all possible cases. This concludes our proof of Theorem 3.1.

## 4. Extracting Safe Regions from Real Sensor Data

The main practical problem with the theoretical results of the previous section is that the sensor output is usually a list of points, not a list of curves. Therefore, a preprocessing stage is needed to convert the raw data into the output list $\Pi$.

Let $L$ be the list of points acquired by the sensor at $q$. $L$ is transformed into a collection $\Pi$ of polygonal lines called polylines. The polyline extraction algorithm operates in two steps: (1) group data into clusters, and (2) fit a polyline to each cluster. The goal of clustering is to group points that can be traced back to the same object in $\mathcal{W}$. Clustering is done using thresholds selected according to the sensor's accuracy.

The points in each cluster are fitted with a polyline so that every data point lies within a distance $\epsilon$ from a line segment, while minimizing the number of vertices in the polyline. The computation takes advantage of the fact that the data delivered by polar sensors satisfy an ordering constraint along the noise-free $\theta$-coordinate axis. By applying the mapping $u=\cos \theta / \sin \theta, v=1 /(r \sin \theta)$, the problem is transformed into a linear fit of the form $v=a+b u$ (which maps to $b x+a y=1$ in Cartesian $(x, y)$-space). Several algorithms exist to find polylines in $(u, v)$-space. We used a divide-and-conquer algorithm. Examples of our polyline-fit technique with real sensor data can be found in [5].

### 4.1. Corners

Corners pose a problem even under idealized conditions. Suppose the robot is surrounded by one or several wedge-shaped walls oriented toward the sensor. The sensor is then unable to see any of these wedges, and the safe region is empty. This is not a failure of our mathematical analysis, but a physical limitation of the sensor. This limitation was not taken into account by Definition 2.1, along with several others (e.g., that some surfaces could be perfect mirrors). We can only assume that the angle between any pair of incident walls is large enough such that at least one section at either side of the corner is visible to the sensor. Or that the corner itself is not sharp enough to remain undetected by the sensor (i.e., the corner has "thickness").

Under the above assumptions, we generalize the concept of a surface normal to include corners. The normal $\boldsymbol{n}$ to $\partial \mathcal{W}$ at a corner is the average of $\boldsymbol{n}^{+}$and $\boldsymbol{n}^{-}$, where $\boldsymbol{n}^{+}$and $\boldsymbol{n}^{-}$are the normals to $\partial \mathcal{W}$ immediately after and before the corner (Figure $3(a)$ ). The corner is visible if the conditions of Definition 2.1 are satisfied for this generalized $\boldsymbol{n}$. That is, a corner behaves like any other point in $\partial \mathcal{W}$, as long as our hypotheses about $\partial \mathcal{W}$ hold true.

The system described in [5] expects corners to have thickness, and therefore to be detectable by the sensor. Under this supposition, it is easy to verify that any wedgeshaped object within range is visible if $\tau \geq 45 \mathrm{deg}$ (Figure 3(b)).

## 5. A Next-Best View Algorithm

In a static environment, a safe region remains safe under the union operation. Hence, the layout model can be expanded iteratively. A first partial layout - a local safe region - is constructed from the data acquired by the range sensor at the robot's initial position $q_{0}$. At each iteration, the algorithm updates the layout model by computing the union of the safe region build so far with the local safe region generated at the new position $q_{k}$. The new safe region is then used to select the next sensing position $q_{k+1}$. To compute this next-best-view position, the procedure first generates a set of potential candidates. Next, it evaluates each candidate according to both the expected gain of information that will be sensed at this position, and the motion cost required to move there. These steps are described below.


Figure 4. Next-best view computation after 5 sensing operations have already taken place: (a) candidate generation; $(b)$ evaluation of one candidate.

### 5.1. Model Alignment and Merging

Let $\left\langle\Pi_{g}\left(q_{k-1}\right), S_{g}\left(q_{k-1}\right)\right\rangle$ be the partial global model built at $q_{k-1}$. The term $S_{g}\left(q_{k-1}\right)$ is the union of all local safe regions up to stage $k-1$. The boundary of $S_{g}\left(q_{k-1}\right)$ is composed of free and solid curves, the latter representing physical sections of $\partial \mathcal{W}$. Let $\Pi_{g}\left(q_{k-1}\right)$ be the list of solid curves in the boundary of $S_{g}\left(q_{k-1}\right)$.

The robot performs a sensing operation once it moves into a new location $q_{k}$. From the local measurement $\Pi_{l}\left(q_{k}\right)$, we compute a local safe region $s_{l}\left(q_{k}\right)$ using the techniques from Section 3. Let $\left\langle\Pi_{l}\left(q_{k}\right), s_{l}\left(q_{k}\right)\right\rangle$ be the local model at $q_{k}$.

Suppose there exists an algorithm ALIGN that computes the transformation $T$ aligning the line segments in $\Pi_{l}\left(q_{k}\right)$ with those in $\Pi_{g}\left(q_{k-1}\right)$. We will not assume that this technique is perfect: ALIGN computes a correct $T$ only when there is enough overlap between $\Pi_{l}\left(q_{k}\right)$ and $\Pi_{g}\left(q_{k-1}\right)$.

Once $T$ is calculated, the new global safe region $S_{g}\left(q_{k}\right)$ is computed as the union of $T\left(S_{g}\left(q_{k-1}\right)\right)$ and $s_{l}\left(q_{k}\right)$. The new model $\left\langle\Pi_{g}\left(q_{k}\right), S_{g}\left(q_{k}\right)\right\rangle$ is represented in a coordinate frame centered over the robot at its current position $q_{k}$.

### 5.2. Candidate Generation

The next location $q_{k+1}$ must be contained inside $S_{g}\left(q_{k}\right)$. Otherwise, $q_{k+1}$ would not be reachable from the current position $q_{k} \cdot{ }^{1}$

We generate at random a number of possible NBV candidates in $S_{g}\left(q_{k}\right)$ within the vicinity of the free curves bounding $S_{g}\left(q_{k}\right)$ (Figure $4(a)$ ). For each possible candidate $q$, we compute the total length $\zeta\left(S_{g}\left(q_{k}\right), q\right)$ of the non-free curves bounding $S_{g}\left(q_{k}\right)$ that are visible from $q$ under Definition 2.1 (this operation is done using a linesweep technique [10]). $\zeta$ is the measure of the expected overlap between a new image $\Pi_{l}(q)$ and the current list of solid curves $\Pi_{g}\left(q_{k}\right)$. If $\zeta\left(S_{g}\left(q_{k}\right), q\right)$ is greater than some threshold, then $q$ is actually selected as an NBV candidate. This filtering stage ensures that the function ALIGN will successfully find a transform $T$.

[^0]
### 5.3. Evaluation of candidates

To decide whether a position $q$ in $S_{g}\left(q_{k}\right)$ is a good candidate for $q_{k+1}$ we must estimate how much new information about the workspace we expect to obtain at $q$ - i.e., $q$ should potentially see large unexplored areas through the free boundary of $S_{g}\left(q_{k}\right)$.

The score of every NBV candidate $q$ is given by the function $g(q)=$ $A(q) \exp \left(-\lambda L\left(q, q_{k}\right)\right)$, where $\lambda$ is a positive constant, $L\left(q, q_{k}\right)$ is the length of the shortest path connecting $q_{k}$ with $q$, and $A(q)$ is a measure of the unexplored area of the environment that may be visible from $q$ (see next paragraph). $q_{k+1}$ is selected as the sample $q$ that maximizes the function $g(q)$. The factor $\lambda$ weights the relative cost of motion with respect to visibility gains. $\lambda=0$ implies that the map builder incurs no cost while moving, and the NBV planner is allowed to select new locations exclusively in terms of their potential visibility gain. $\lambda \gg 0$ implies that motion is so costly that locations close to $q_{k}$ are preferred over distant ones, as long as they produce a marginal gain in visibility.
Computation of $\mathbf{A}(\mathbf{q})$ We measure the potential visibility gain of each candidate $q$ as a function of the area $A(q)$ outside the current safe region that may be visible through the free curves bounding $S_{g}\left(q_{k}\right)$ (Figure $4(b)$ ). For polygonal models, $A(q)$ can be computed by the same ray-sweep algorithm used to compute classic visibility regions [10], with the following modifications:

1. The sweeping ray may cross an arbitrary number of free edges before hitting a solid one. Therefore, the computation-time of the ray-sweep algorithm becomes $\mathrm{O}\left(n \log (n)+n k_{\mathrm{f}}\right)$, where $k_{\mathrm{f}}$ is the number of free edges bounding $S_{g}\left(q_{k}\right)$.
2. The resultant visible region is cropped to satisfy the range restrictions of the sensor. This operation can be done in $\mathrm{O}\left(n k_{\mathrm{f}}\right)$.

### 5.4. Termination Condition

If $S_{g}\left(q_{k}\right)$ contains no free curves, the 2-D layout is assumed to be complete; otherwise, $S_{g}\left(q_{k}\right)$ is passed to the next iteration of the mapping process. A weaker termination test is employed in practice: the length of any remaining free curve is smaller than a specified threshold.

### 5.5. Iterative Next-Best View Algorithm

The iterative NBV algorithm is summarized below:
Algorithm Iterative Next-Best View
Input: A new sensing position $q_{k}$ and the local measurement $\Pi_{l}\left(q_{k}\right)$ An image alignment function $T=\operatorname{ALIGN}\left(\Pi_{l}\left(q_{k}\right), \Pi_{g}\left(q_{k-1}\right)\right)$
The number of samples $m$, and a weighting constant $\lambda>0$
Output: A next-best view position $q_{k+1}$
Compute the local safe region $s_{l}(q k)$. Set the list of samples $\mathcal{N}_{\text {sam }}=\emptyset$.
Compute $T=\operatorname{ALIGN}\left(\Pi_{l}\left(q_{k}\right), \Pi_{g}\left(q_{k-1}\right)\right)$, and the union $S_{g}\left(q_{k}\right)=$ $s_{l}\left(q_{k}\right) \bigcup T\left(S_{g}\left(q_{k-1}\right)\right)$.
3. Repeat until the size of $\mathcal{N}_{\text {sam }}$ is greater or equal than $m$ :
(a) Randomly generate $q \in S_{g}\left(q_{k}\right)$ in the vicinity the free curves bounding $S_{g}\left(q_{k}\right)$.
(b) If $\zeta\left(S_{g}\left(q_{k}\right), q\right)$ is below the requirements of ALIGN, discard $q$ and repeat Step 3.
(c) Compute $A(q)$ and $L\left(q, q_{k}\right)$. Add $q$ to $\mathcal{N}_{\text {sam }}$ and repeat Step 3.
4. Select $q_{k+1} \in \mathcal{N}_{\text {sam }}$ maximizing $A(q) \exp \left(-\lambda L\left(q, q_{k}\right)\right)$ as the next-best view.


Figure 5. A run around the Robotics Lab. at Stanford University.

## 6. Experiments

The map-building system was implemented on a Nomadic SuperScout robot. The on-board computer is a Pentium 233 MMX, connected to the local-area network via 2 Mbs radio-Ethernet. The robot is equipped with a laser range sensor from Sick Optic Electronic which uses a time-of-flight technique to measure distances. The NBV planner runs off-board in a Pentium II 450 MHz Dell computer. The software was written in C++ and uses geometric functions from the LEDA library [8].

The sensor acquires 360 points in a single 180-deg scan request. A 360-deg view is obtained by taking 3 scans. The sensor readings where observed to be reliable within a range of 6.5 mts , at grazing angles not exceeding $\tau=85 \mathrm{deg}$. For the NBV planner, $\lambda=20 \mathrm{~cm}^{-1}$, a value that prevents the robot from oscillating back and forth between regions with similar visibility gains.

An experimental run is shown in Figure 5. The robot mapped a section of the Robotics Lab. at Stanford U. The first 6 iterations are shown in (a). At the corridor intersection, the robot faces three choices, including going into an office. Nevertheless, the planner opted to continue moving along a corridor, all the way into the upper hall $(b)$. Glass is transparent to the sensor's laser, so the robot failed to detect the glass door indicated in (b). At this point, the operator overrode the decision of the NBV planner, who interpreted the vicinity of the glass door as the threshold of an unexplored open area. Finally, in (c), the robot moved down the second hall until it reached the lab's lounge. The planner decided then to send the robot to explore this newly detected area.

## 7. Conclusion

Motion planning for model building applications has received little attention so far despite its potential to improve the efficiency of autonomous mapping. In this paper we introduced the concept of safe region, and described how it can be used to produce collision-free motions and next-best view locations under image-alignment considerations. Our research combines theoretical investigation of planning problems with simplified visibility models to produce algorithms that reach a compromise between algorithmic rigor and system practice. The result is a system able to construct models of realistic scenes.
Acknowledgments: This work was funded by DARPA/Army contract DAAE07-98L027, ARO MURI grant DAAH04-96-1-007, and NSF grant IIS-9619625.

## References

[1] R. Chatila and J.P. Laumond. Position referencing and consistent world modeling for mobile robots. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 138-143, 1985.
[2] H. Choset and J. Burdick. Sensor based motion planning: The hierarchical generalized voronoi diagram. In J.-P. Laumond and M. Overmars, editors, Proc. 2nd Workshop on Algorithmic Foundations of Robotics. A.K. Peters, Wellesley, MA, 1996.
[3] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In Proc. ACM SIGGRAPH, 1996.
[4] A. Elfes. Sonar-based real world mapping and navigation. IEEE J. Robotics and Automation, RA-3(3):249-265, 1987.
[5] H. González-Banos, A. Efrat, J.C. Latombe, E. Mao, and T.M. Murali. Planning robot motion strategies for efficient model construction. In Robotics Research - The Eight Int. Symp., Salt Lake City, UT, 1999. Final proceedings to appear.
[6] B. Kuipers, R. Froom, W.K. Lee, and D. Pierce. The semantic hierarchy in robot learning. In J. Connell and S. Mahadevan, editors, Robot Learning. Kluwer Academic Publishers, Boston, MA, 1993.
[7] J.J. Leonard and H.F. Durrant-Whyte. Stochastic multisensory data fusion for mobile robot location and environment modeling. In Proc. IEEE Int. Conf. on Intelligent Robot Syst., 1991.
[8] K. Mehlhorn and St. Nahër. LEDA: A Platform of Combinatorial and Geometric Computing. Cambridge University Press, Cambridge, UK, 1999.
[9] P. Moutarlier and R. Chatila. Stochastic multisensory data fusion for mobile robot location and environment modeling. In H. Miura and S. Arimoto, editors, Robotics Research - The 5th Int. Symp., pages 85-94. MIT Press, Cambridge, MA, 1989.
[10] J. O'Rourke. Visibility. In J.E. Goodman and J. O'Rourke, editors, Handbook of Discrete and Computational Geometry, pages 467-479. CRC Press, Boca Raton, FL, 1997.
[11] R. Pito. A sensor based solution to the next best view problem. In Proc. IEEE 13th Int. Conf. on Pattern Recognition, volume 1, pages 941-5, 1996.
[12] S. Teller. Automated urban model acquisition: Project rationale and status. In Proc. 1998 DARPA Image Understanding Workshop. DARPA, 1998.
[13] G. Turk and M. Levoy. Zippered polygon meshes from range images. In Proc. ACM SIGGRAPH, pages 311-318, 1994.


[^0]:    ${ }^{1}$ Strictly speaking, $S_{g}\left(q_{k}\right)$ must first be shrunk by the radius of the robot before computing a safe route.

