

A family of skeletons for motion planning and geometric reasoning applications

ATA A. EFTEKHARIAN AND HOREA T. ILIEȘ

Department of Mechanical Engineering, University of Connecticut, Storrs, Connecticut, USA

(RECEIVED September 20, 2010; ACCEPTED February 7, 2011)

Abstract

The task of planning a path between two spatial configurations of an artifact moving among obstacles is an important problem in practically all geometrically intensive applications. Despite the ubiquity of the problem, the existing approaches make specific limiting assumptions about the geometry and mobility of the obstacles, or those of the environment in which the motion of the artifact takes place. We present a strategy to construct a family of paths, or roadmaps, for two- and three-dimensional solids moving in an evolving environment that can undergo drastic topological changes. Our approach is based on a potent paradigm for constructing geometric skeletons that relies on constructive representations of shapes with R-functions that operate on real-valued half-spaces as logic operations. We describe a family of skeletons that have the same homotopy as that of the environment and contains the medial axis as a special case. Of importance, our skeletons can be designed so that they are “attracted to” or “repulsed by” prescribed spatial sites of the environment. Moreover, the R-function formulation suggests the new concept of a medial zone, which can be thought of as a “thick” skeleton with significant applications for motion planning and other geometric reasoning applications. Our approach can handle problems in which the environment is not fully known a priori, and intrinsically supports local and parallel skeleton computations for domains with rigid or evolving boundaries. Furthermore, our path planning algorithm can be implemented in any commercial geometric kernel, and has attractive computational properties. The capability of the proposed technique are explored through several examples designed to simulate highly dynamic environments.

Keywords: Boolean Expression; Half-Space; Medial Zone; Path Planning; R-Function

1. INTRODUCTION

Motion planning is ubiquitous in many geometrically intensive applications involving moving objects such as robotics, virtual reality, computer graphics, computer-aided design, and manufacturing, computer-aided surgery, computational biology, and many others. It heavily relies on spatial information about the environment in which the navigation takes place, and typically focuses on generating a “useful” motion that takes a moving object, a robot, or a vehicle, between two prescribed configurations subject to some constraints such as minimum length of the path, collision-free trajectories, and intermediate sites.¹ One of the most useful and common abstractions of this problem is to transform the motion planning

of an object in the Euclidean space into the path planning of a point in the configuration space (or C-space) of the object (Latombe, 1991).

Consider a moving object leaving from an initial position and moving in the Euclidean space E^d ($d = 2, 3$) toward a prescribed final destination. Such a vehicle can be a mobile robot, an autonomous vehicle, a camera flying through a scene, or collection of spatial entities,² a mechanical part that is being assembled or transported, and so forth. The object may be constrained by intermediate sites that must be either approached or avoided during the motion, such as a fueling/loading/unloading station or a sightseeing location. The environment itself can be static or dynamic, case in which the obstacles themselves or the boundary of the environment are changing during the motion. The task is to find a path from start to finish that satisfies the intermediate site constraints and is subjected to other specific constraints as illustrated in Figure 1. In its most general form, finding a path in such an

Reprint requests to: Horea T. Ilieș, Department of Mechanical Engineering, 366 United Technologies Building, 191 Auditorium Road, Unit 3139, University of Connecticut, Storrs, CT 06269-3139, USA. E-mail: ilies@enr.uconn.edu

¹ Motion planning has been recently extended to include differential constraints that restrict allowable velocities during the motion, and use sampling-based planning as discussed, for example, in LaValle (2006).

² See Snavely et al. (2008) for a very interesting application.




Fig. 1. An object moving in an evolving environment with one intermediate site that must be approached during the motion. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

environment is a hard problem. Consequently, the existing approaches address simplified versions of this problem.

A large amount of research has been devoted to motion planning in static environments with both exact and approximate methods. The existing approaches for finding the motion, which cannot be easily extended to dynamic environments, can be classified into the following:

1. *roadmap methods* that plan curves in either the configuration space of the moving object or in the Euclidean space in which the object moves. There has been some success in constructing three-dimensional (3-D) configuration spaces (see, e.g., Sacks et al., 1999), but computing C-spaces explicitly remains a difficult problem, which is exponential in the number of dimensions of the C-space (Latombe, 1991). Consequently, computing roadmaps in the Euclidean space has remained a strong area of research. Current motion planning algorithms that compute roadmaps rely on visibility graphs (Jiang et al., 1999; Neus & Maouche, 2005), Voronoi diagrams (Diaz de León & Sossa, 1998; Wilmarth et al., 1999; Foskey et al., 2001; Garber & Lin, 2002; Lee & Choset, 2005; Geraerts & Overmars, 2007), silhouette curves of semialgebraic sets (Canny, 1993), or probabilistic roadmap planners (Kavraki et al., 1996);
2. *potential field methods* that introduce artificial attractive/repulsive field in the environment can conceptually handle both static and moving obstacles (Rimon & Koditschek, 1992; Xidias et al., 2007). These methods cannot take into consideration the exact geometry of the obstacles, and therefore they cannot be applied to those problems in which the moving object is moving in close proximity with the obstacles. It is of more importance that potential field methods introduce local minima of the potential field that can “trap” the moving object;
3. *approaches based on cell decompositions*, which are inherently approximate. These methods perform a disjoint decomposition of the free space, construct the connectivity graph connecting these cells, and perform

graph search algorithms to select a collision free path (Choset, 2000; Lingelbach, 2004).

An approximately time-optimal trajectory is proposed in van den Berg and Overmars (2005) that uses a precomputed roadmap for the static part of the scene to determine the collision-free trajectory among the moving obstacles. This approach requires the complete environment and motions to be known ahead of time. Other approaches that incorporate into the formulation time as the extra dimension have been discussed in Ishikawa (1991) and Fraichard (1993). More general approaches to dynamic motion planning couple a precomputed probabilistic road map and a cell decomposition to dynamically determine which cells of the roadmap are affected by the motion of the obstacles (Kallman & Mataric, 2004). These methods produce successful paths when the dynamic subset of the environment is much smaller than the whole environment, and when changes to the precomputed roadmap are fairly minor. Motion planning in dynamic environments in which both moving obstacles and target are moving with known velocities has been addressed in Masehian and Katebi (2007). Because of the assumption of known velocity vectors, this approach cannot handle the case when the number, type, and motion of the moving obstacles are not known in advance.

A fast motion planning algorithm in a two-dimensional (2-D) dynamic environment based on generalized Voronoi diagrams has been described in Hoff et al. (2000). The Voronoi decomposition is computed for the prescribed “sites”—the obstacles of the domain. For each pixel in the domain they compute the closest site and shortest distance to that site. All pixels closest to the same site get assigned the same color so that different Voronoi cells will have different colors. The algorithm has been implemented on the graphics processing unit, which allows the computation of the Voronoi decomposition at each time step.

1.1. Main contributions of our work

In this paper we present a strategy to develop a family of paths, or roadmaps, for 2-D and 3-D solids³ moving in highly dynamic environments that can undergo drastic topological changes. The main contribution of this paper is the definition of a class of skeletons that have the same homotopy as that of the environment and contains the medial axis as a special case. It is important that our skeletons can be designed so that they are attracted to or repulsed by prescribed sites in the environment, whereas the attraction and repulsion are controlled by a set of degrees of freedom (see Section 2.6). This can effectively alter the planned path to accommodate specific geometric constraints such as approaching a refueling station or a sightseeing location. Furthermore, we exploit the concept of a medial zone that we introduced in Eftekharian

³ These solids are closed, bounded, regular, and semianalytic sets (Requicha, 1980).

and Ilieş (2010) to control the subset of the environment in which the roadmaps are being computed, and we show that, for the same domain, the resulting roadmaps are shorter and smoother than those obtained based solely on skeletons.

Once we construct the corresponding skeletons and medial zones, we use the established Dijkstra's algorithm to compute shortest paths of moving points along the skeletons or within the medial zones. Furthermore, the properties of R-functions afford a conservative motion planning of solid objects within the same formulation via level sets. The approach presented in this paper can handle problems in which one or more objects move in environments that are not fully known in advance, and intrinsically supports local and parallel computations of skeleton and medial zones (and, hence, path planning) for domains with rigid or evolving boundaries. Our approach can be implemented in any commercial geometric kernel, and has attractive computational properties.

2. PROBLEM FORMULATION

In this paper we focus on evolving spatial environments with obstacles that can move as well as merge/split. We assume that the initial and final position (i.e., points in the environment) of the path to be planned are known, and that two lists of obstacles that contain the "attractive" obstacles \mathcal{A} as well as "repulsive" obstacles \mathcal{R} are provided along with positive parameters $\lambda \in \mathbb{R}^+$ that quantify the extent by which each obstacle will attract or repulse the path to be computed (see Section 2.6). Each site that must be approached is specified by the user. The task that we investigate in this paper is to compute the shortest path along the skeleton or within the medial zones that connects the prescribed initial and final positions subject to the intermediate constraints, or output a message that one does not exist for the prescribed environment.

2.1. Summary of the approach

Our approach is illustrated in Figure 2 and consists of five, not necessarily sequential, stages:

1. Construct an exact or approximate distance function over the given semianalytic domain.
2. Adjust the distance function according to the "attractive" and "repulsive" factors for obstacles prescribed in \mathcal{A} and \mathcal{R} .
3. Extract the skeleton of the domain as the subset of points where the continuous distance function is non-differentiable. Note that our construction provides an explicit mapping between each half-space defining the environment and each branch of the skeleton, which can play a critical role in local skeleton computations.
4. Alternatively, compute the medial zones as "thick" versions of the skeletons;
5. Compute the shortest collision free path along the skeleton or inside the medial zones subject to the intermediate sites that must be visited.

Steps 1 and 3 above are summarized in Section 2.2 (see also Eftekharian & Ilieş, 2009); step 2 is discussed in Section 2.3, whereas the medial zones are summarized in Section 3 and described in detail in Eftekharian and Ilieş (2010). Section 4 illustrates the usefulness of our approach for several domains that simulate highly dynamic, topologically evolving domains.

2.2. Preliminaries

Geometric skeletons are fundamental concepts in practically all geometrically intensive areas of science and engineering, such as automated finite element meshing, shape manipulation, recognition, and comparison, dimensional reduction in design and analysis, robotic surgery, and a variety of path and motion planning in commercial and defense applications. More recently, skeletons have been used to explore the fundamental geometric problems of folding and unfolding that are the abstraction of some of the most important open problems in science today, such as protein folding, as well as packing and sheet metal bending. There are multiple ways to define a skeleton of a given set, and a variety of definitions have been proposed for different applications. Introduced by Blum (1967) as a tool for image analysis, the medial axis has become one of the mainstream geometric concepts due to the fact that it provides a compact representation of the geometric features of a shape and its topology. The medial axis captures the connectivity of the shape, has a lower dimension than the space itself, and is closely related to the distance function constructed over the same domain.

The concept of medial axis has been described with the help of the fire grass concept (see, e.g., Attali et al., 2008) as follows: if a fire starts from all points of a planar curve at the same time and moves with constant velocities in all directions in the same plane, then the medial axis is the locus of points where the fire (in fact, a moving front) meets itself. The concept can be extended to k -dimensional geometric shapes in \mathbb{R}^k , case in which the medial axis becomes a set of dimension $k - 1$. Intuitively, the points on the medial axis are equally distant to at least two points of the boundary of the domain. Since its formulation, the medial axis has been used as alternative solid modeling representation (Shaham et al., 2004), as well as in many other applications such as shape matching and reconstruction (Liu et al., 1998; Siddiqi et al., 1999; Sebastian et al., 2004; Damon, 2005; van Eede et al., 2006; Goh, 2008), dimensional reduction in boundary value problems (Suresh, 2003), representation and classification of 2-D shapes (Sherbrooke et al., 1995; Pizer et al., 2003; Shah, 2005), human vision (Kimia & Tamrakar, 2002), pattern analysis and shape recognition (Blum & Nagel, 1978; Bookstein, 1979), mesh generation (Sampl, 2000, 2001; Quadros et al., 2004), and so forth.

The mathematical properties of medial axis are fairly well understood (Choi et al., 1997; Chazal and Soufflet, 2004; Attali et al., 2008). However, the practical uses of these skeletons are limited by their notorious computational diffi-




Fig. 2. (a) The three-dimensional distance function of the domain constructed with R-functions, (b) the medial axis of the domain and the shortest path that follows the medial axis of the domain, and (c) the modified skeleton attracted to the intermediate target and the resulting shortest path. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

culties, and their instability under small perturbations of the shape. Although it is possible, in principle, to compute the medial axis exactly for general semialgebraic/analytic sets, we do not seem to have any algorithms for doing so. The most advanced algorithms make significant simplifying assumptions on the geometry of these sets by focusing on planar or piecewise linear shapes. Algebraic planar curve segments whose bisectors admit rational parametrizations are examined in Ramamurthy and Farouki (1999), whereas exact computation of medial axis for polyhedra is described in Culver et al. (2004). The reasons behind these restrictions become apparent once the algebraic difficulties in computing medial axis (MA) are examined (Attali et al., 2008).

Consequently, these difficulties promoted a variety of algorithms that approximate the complex shape by a set bounded by a piecewise linear boundary for which the medial axis can be computed exactly, followed by the extraction of the medial axis of the approximating shape—the so-called pruning step. The computed medial axis has to be postprocessed in order to eliminate the branches that appear in the medial axis due to the shape approximation. The main approaches to approximately compute the medial axis for piecewise linear shapes rely on: Voronoi/Delaunay decompositions of space (Brandt & Algazi, 1992; Amenta et al., 2001; Dey et al., 2003) where the medial axis is the Voronoi graph defined by a piecewise linear approximation of the shape (Attali et al., 2008); solutions of partial differential equations (such as diffusion or Hamilton–Jacobi equations; Siddiqi et al., 2002; Du & Qin, 2004); and level set methods (Kimmel et al., 1995; Gomes & Faugeras, 2000). A recent method to compute the medial axis for curved planar sets (Cao & Liu, 2008) is iteratively tracing the Frenet frames for pairs of boundary curves that bound a closed planar domain. This algorithm does not generalize to 3-D space and its stability appears to be problematic in 2-D due to its iterative marching along the boundary curves.

2.2.1. R-functions as logic operators on real-valued half-spaces

For any closed subset Ω of \mathbb{E}^d , one can construct a C^∞ function that vanishes on the boundary $\partial\Omega$ of Ω . It is known

that such shapes are in fact semianalytic sets of points that can be constructed as a finite Boolean combination of real analytic functions $f_i \geq 0$. This in turn suggests an approach to construct a C^n function over a semianalytic subset Ω of \mathbb{E}^d by subdividing the boundary of Ω in primitive half-spaces f_i , followed by a combination of f_i into a single predicate using the standard Boolean logic operators AND, OR, or NOT (Shapiro, 2007).

R-functions have been invented in the 1960s by V.L. Rvachev, who called these functions “logically charged functions.” These functions provide the means to construct a C^n function over a domain defined by primitive half-spaces. The main contribution of the theory of R-functions to the topic of this paper is to replace these logical operations by real-valued functions, which generates an implicit representation for any semianalytic set Ω . One important feature of these real valued functions is that their sign is completely determined by the sign of their arguments, and is independent of any of their magnitudes.

There are many systems of sufficiently complete R-functions. One such system is known as principal system of R-functions

$$R_\alpha(\Delta): \frac{1}{1 + \alpha} \left(x_1 + x_2 \pm \sqrt{x_1^2 + x_2^2 - 2\alpha x_1 x_2} \right), \quad (1)$$

where (+) and (−) signs correspond to the R-conjunction ($x_1 \vee_\alpha x_2$) and R-disjunction ($x_1 \wedge_\alpha x_2$) respectively of two real variables x_1 and x_2 . By varying the value of α , we obtain different systems of R-functions. In particular, by setting $\alpha = 1$ in Eq. (1) we obtain the $R_1(\Delta)$ system of R-functions, whereas a value $\alpha = 0$ in Eq. (1) would result in the $R_0(\Delta)$ system (see Fig. 3).

Of importance, when expression under the square root in Eq. (1) becomes zero, the corresponding implicit function becomes nondifferentiable. Because the medial axis of a semianalytic domain corresponds to the points where the distance function is nondifferentiable, we will be looking for points of the implicit function where the expression under the square root vanishes.




Fig. 3. Three implicit representations of a polygonal domain that correspond to (a) $\alpha = 0$, (b) $\alpha = 0.5$, and (c) $\alpha = 1$. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

Figure 3 shows the difference between the implicit functions constructed over a polygonal domain with the $R_\alpha(\Delta)$ system for three values of α . One can see that for $\alpha \neq 1$, the resulting functions are differentiable (except at the origin, not shown). In contrast, the implicit function corresponding to $R_1(\Delta)$ shown in Figure 3c is made of piecewise planar patches that intersect at piecewise linear edges. In fact, by projecting the edges of the R-function surface on the plane of the polygon we obtain a skeleton of the polygon. However, how do we construct the function, and which skeleton do we obtain?

2.2.2. Constructing Boolean expressions

Conceptually, the problem of constructing a Boolean expression for a domain bounded by half-spaces is the same as that of converting a boundary representation (B-rep) into a constructive solid geometry (CSG) representation in solid modeling.

Several efficient algorithms for performing the B-rep to CSG conversion for polygons are known (Tor & Middle-ditch, 1984; Dobkin et al., 1988). However, the complexity of the problem quickly escalates with the increase in the complexity of the boundary of the domain. Separating the boundary into primitive pieces is no longer sufficient to construct the Boolean expression of the domain, because additional half-spaces need to be introduced, which are the so-called separating half-spaces. Determining a sufficient set of separating half-spaces is the critical step of any such algorithm, but the problem is not well understood in general. Solutions exist for planar domains bounded by linear and curved edges that are subsets of convex curves (Shapiro, 2001) and solids bounded by linear or quadric surfaces (Shapiro & Vossler, 1993).

Note that the problem of B-rep to CSG conversion is well defined by using the natural half-spaces of the domain, but a CSG expression may not exist for a given set of half-spaces. This necessary and sufficient condition for the existence of a CSG expression is formally described by the “describability theorem” in Shapiro and Vossler (1991). However, if a canonical CSG expression exists for a given set of half-spaces bounding a specific domain, then this canonical CSG expression is unique. Observe that we do not require the canonical CSG expression, because all CSG expressions that are valid for a given domain will describe the same set of points.

Planar domains. The Boolean set representation of a polygonal domain can be computed based on the convex deficiency tree (Dobkin et al., 1988), which treats each polygon as its convex hull minus a finite number of concavities. Note that the polygon is a closed set, so the subtraction of concavities must necessarily be regularized (i.e., one must use regularized Boolean operations). Finally, we perform a syntactic substitution to replace the union and intersection with the R-disjunction and R-conjunction given in Eq. (1), which results in an R-function expression whose zero set is the original planar domain. By following this procedure, we obtain an

R-function expression that corresponds to the exact distance function for any convex planar domain polygon, and an approximate distance function for a concave domain. Note that such a distance function is obtained from the principal system of R-functions given in Eq. (1) by setting the value of $\alpha = 1$, that is, the $R_1(\Delta)$ system. Values of $\alpha < 1$ correspond to implicit functions over the same domain that have established differential properties (Shapiro, 2007). These approximate distance functions can be converted into an exact distance function by introducing additional half-spaces at the concave vertices of the domain (Eftekharian & Ilieş, 2009). Specifically, each concave vertex requires a conical half-space with a half angle of $\pi/4$ and two separating/trimming half-spaces that are normal to the boundary curves that are incident at the concave vertex.

3-D domains. It is important to note that the construction algorithms for simple polygons can be extended to some other point sets, such as curved polygons (Shapiro, 2001), 3-D polyhedra, and more general 3-D solids (Shapiro, 1991; Buchele & Crawford, 2003). However, converting the resulting approximate distance function into an “exact” distance function for such domains requires the computation of additional halfspaces as discussed above, which can only be computed approximately in many situations. The alternative that we take in this paper is to compute distance functions to each face bounding the 3-D domain, and combine them with the R-disjunction as described in more detail in Eftekharian and Ilieş (2010).

2.2.3. Adding/removing obstacles

Obstacles in the environment can be represented as holes/voids of the domain. One of the main advantages of our approach to compute skeletons based on R-functions is that obstacles can be easily added to the R-function expression representing Ω (see also Eftekharian & Ilieş, 2009). The R-function expression corresponding to the obstacle is subtracted from the R-function expression defining the boundary of Ω according to the R-subtraction⁴

$$R_1(\Delta) = \frac{1}{1+\alpha} \left(f_1 - f_2 - \sqrt{f_1^2 + f_2^2 + 2f_1 f_2} \right), \quad (2)$$

where f_1 is the function describing the outer domain and f_2 represents the hole. Figure 4 shows how obstacles can be added to the main space. Figure 4a–c illustrates this process in which two obstacles are sequentially added to the environment, then one of them is rotated. The corresponding distance functions are illustrated in Figure 4d–f.

One important property of the R-functions is that their sign is independent of the magnitude of each half-space involved in the R-function expression, and depends only on the sign of

⁴ Recall that the subtraction of two point sets A and B is usually defined in terms of the Boolean operations as $A - B = A \cap B^c$, where B^c is the complement of B which implies that for solids one must use regularized set operations.




Fig. 4. Obstacles can be added to (or removed from) the environment. (a)–(c) Distance functions to individual obstacles are illustrated, and (d)–(f) the resulting distance functions of the domain are shown. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

each half-space. Consequently, holes can be easily relocated and resized to perform, for example, parametric or topology optimization (Chen et al., 2007; Luo et al., 2007). These modifications exploit the connection between implicit functions defined with R-functions and level sets so that the boundary of the domain defined by the R-function expression is the zero level set of the corresponding implicit function, and moving boundaries of Ω can be handled as easily (Eftekharian and Ilies, 2009).

2.3. From distance functions to collision-free trajectories

Constructing the R-function expression described above results in the distance function of any 2-D or 3-D bounded, regular, closed and semianalytic domain. The ridges of the distance function contain the points where the distance function is nondifferentiable, and therefore, belong to the medial axis of the domain.

In Eftekharian and Ilies (2009) we explored two strategies for extracting the ridges of the distance function. Specifically,

1. for the planar case, if the distance functions H_i corresponding to each individual half-space f_i are constructed in a commercial geometric kernel as a standard 3-D NURBS surface, then the distance function corresponding to domain Ω can be obtained by combining the individual H_i according to the Boolean expression developed in the previous step. This outputs individual curve segments/branches of the medial axis.
2. Alternatively, one can take advantage of the fact that the distance function is not differentiable at the ridge points. If one assumes that the first derivative of a given continuous function is large (it reaches a local extremum) wherever the second derivative has a zero crossing, one can use Laplacian based methods for ridge de-

tection, which are widely used in image analysis and computer vision. This leads to a numerical procedure that outputs points of the medial axis, which needs to be followed by a segmentation of these points into branches of the medial axis. One practical way to perform the medial axis segmentation (or branching) is to determine the half-spaces H_i that evaluate to zero for each point of the medial axis. The advantage of this approach is that it can be applied to both 2-D and 3-D path planning problems.

Regardless of the ridge extraction algorithm being employed, by constructing a single function describing the boundary of each obstacle we can significantly reduce the number of branches of the medial axis, which will significantly reduce the computational cost of the path planning algorithm. This is illustrated in Figure 5, which shows a polygonal domain and its exact distance function in Figure 5a. The rectangular obstacle is bounded by four half-spaces H_1, \dots, H_4 , to which we add trimmed conical half-spaces C_1, \dots, C_4 . These half-spaces can be treated individually or collectively during the segmentation, as mentioned above. Note that even though we use linear half-spaces for illustrative purposes, the approach remains valid for all semianalytic half-spaces.

If each individual half-space bounding the obstacle is treated individually during the segmentation, then one half-space defining the outer boundary of the environment, say H_5 , will generate four separate branches of the medial axis as shown in Figure 5b. Alternatively, if we construct one single Boolean expression H_0 for the boundary of the obstacle as

$$H_0 = (H_1 \cup H_2 \cup H_3 \cup H_4) \cup (C_1 \cup C_2 \cup C_3 \cup C_4), \quad (3)$$

then the segmentation will effectively combine the four individual branches of the medial axis into one single branch (see Fig. 5c). Repeating this procedure for the remaining outer




Fig. 5. Combining the half-spaces defining the obstacles into a single function reduces the number of branches of the medial axis. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

boundary of Ω will generate the remaining branches of the medial axis according to

$$\Omega = [(H_1 \cup H_2 \cup H_3 \cup H_4) \cup (C_1 \cup C_2 \cup C_3 \cup C_4)] \cap (H_5 \cap H_6 \cap H_7 \cap H_8) \tag{4}$$

and replace the Boolean operations with the R-conjunction and R-disjunction given in Eq. (1).

Figure 6 illustrates the impact of the two approaches on the number of branches of the medial axis for a simple planar domain. In this example, by combining the half-spaces bounding each obstacle as described above, we reduce the number of branches of the medial axis from 30 to 14. The

number of nodes in the graph will therefore be reduced from 32 to 12.

After the medial axis segmentation, the last step in planning a path for a moving point is to perform a graph search algorithm on a weighted graph. Each critical point becomes a vertex in this graph and each branch will be represented as a weighted edge whose weight is proportional to the length of the branch. It should be apparent that the “smaller” the graph the faster the computations. In this work we used the established Dijkstra’s algorithm (Dijkstra, 1959) that has been widely used in motion planning and routing problems due to its simplicity, efficiency, and robustness. However, one can use any of the other existing graph search algorithms to generate collision-free paths in this weighted graph.




Fig. 6. The half-spaces bounding each individual obstacle are combined so that each obstacle will be defined by one single function, which reduces the number of medial axis branches from 30 to 14. (a) Dashed lines are branches created by intersection of conical half-spaces with the outer half-spaces (see Section 2.2.2). The black dots on the medial axis separate the branches of the medial axis. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

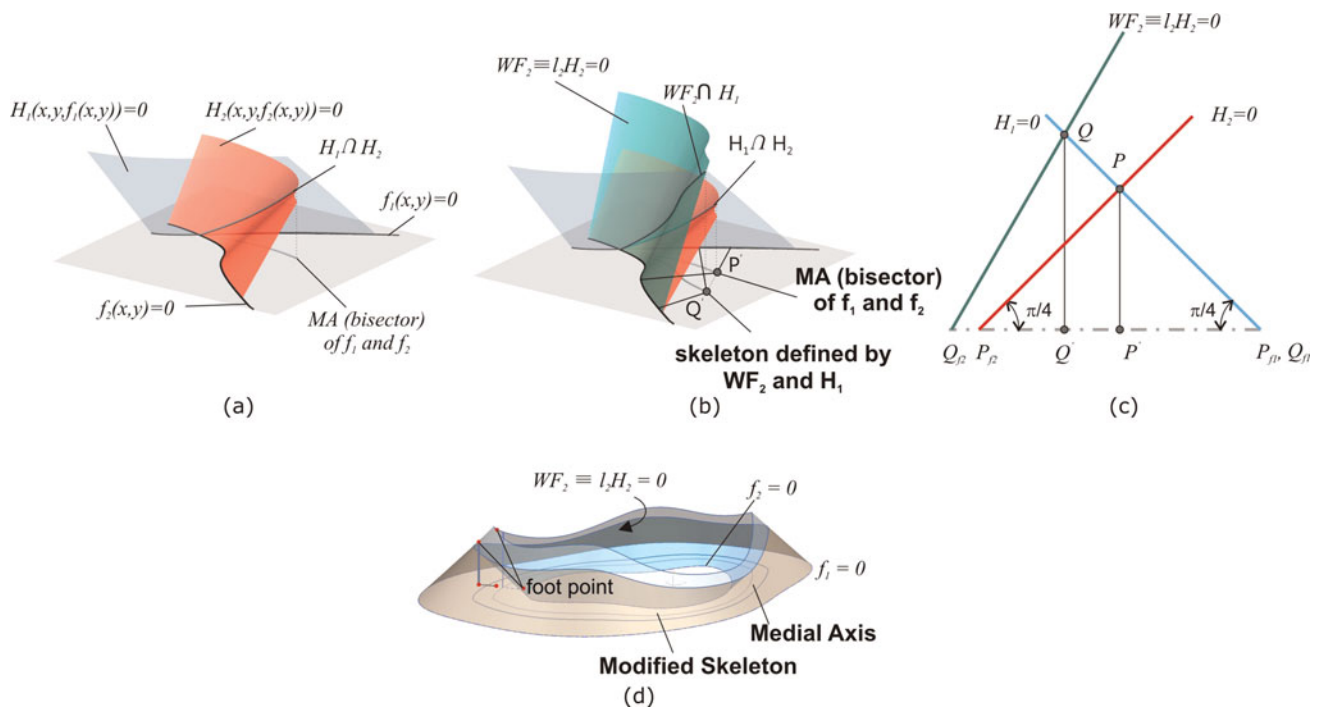


Fig. 7. (a) Exact distance functions $H_1 = 0$ and $H_2 = 0$ for two planar curves defined by $f_1(x, y) = 0$ and $f_2(x, y) = 0$, (b) a weighted distance function $WF_2 = \lambda_2 H_2 = 0$ results in a different skeleton than the medial axis, (c) the triangles formed by $P_i, P'_i (i = 1, 2)$ and the corresponding footpoints flattened on the plane, and (d) a single weight λ is applied on the distance function defining the obstacle $f_2 = 0$ regardless of how many half-spaces contribute to $f_2 = 0$. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

Note that some domains admit more than one shortest path, and that all standard graph search algorithms can be modified to produce all possible shortest paths. The interested reader is referred to Takaoka (2005), as well as to standard textbooks on algorithms such as Dasgupta et al. (2008).

2.4. A family of skeletons that contains the medial axis

To summarize the discussion above, we are first constructing the distance function over the domain from individual half-spaces defining the boundary of the domain (both outer bound-




Fig. 8. Medial zones controlled by the prescribed α values. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

ary and obstacles) as described in Section 2.2.2. As a result, we obtain a continuous piecewise smooth surface, whose zero level set is the boundary of the domain. This construction is followed by the extraction of the medial axis as the subset for which the continuous distance function is nondifferentiable.

However, by altering the distance functions defined by each half-space $f_i \geq 0$, we can deform the medial axis such that the deformed skeleton will be “attracted” or “repulsed” by specific obstacles inside the domain or by specific half-spaces defining the outer boundary of the environment.

In our construction above, we define each halfspace of the original domain as an implicit function $f_i \geq 0$. Each half-space f_i in the d dimensional space \mathbb{E}^d induces a higher dimensional half-space in \mathbb{E}^{d+1} denoted by $H_i \geq 0$.

We will describe the concept with the help of a 2-D example, but this discussion naturally extends to 3-D environments. Consider the example shown in Figure 7, which shows two planar half-spaces f_1 and f_2 , as well as their exact distance functions that are the boundaries of the corresponding 3-D half-spaces defined by equation $H_i = 0$, namely,

$$\partial H_1 = \{P(x, y, f_1(x, y)) \in \mathbb{E}^3 | H_1(x, y, f_1(x, y)) = 0\}, \quad (5)$$

$$\partial H_2 = \{P(x, y, f_2(x, y)) \in \mathbb{E}^3 | H_2(x, y, f_2(x, y)) = 0\}. \quad (6)$$

Assume that the two surfaces ∂H_1 and ∂H_2 are the exact distance functions to curves $f_1 = 0$ and $f_2 = 0$, and that point P is on the intersection curve between ∂H_1 and ∂H_2 as illustrated in Figure 7. By definition, both ∂H_1 and ∂H_2 make with the (x, y) plane an angle of $\pi/4$ at the footpoints of P on f_1 and f_2 , for all points $P \in \partial H_1 \cap \partial H_2$. Denote these footpoints by P_{f_1} and P_{f_2} . Furthermore, assume that $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$ are two positive real numbers that multiply H_1 and H_2 , respectively, such that

$$WF_1 = \lambda_1 H_1(x, y, f(x, y)), \quad WF_2 = \lambda_2 H_2(x, y, f(x, y)).$$

Clearly, the functions WF_1 and WF_2 are distance functions for curves $f_1(x, y) = 0$ and $f_2(x, y) = 0$ if and only if $\lambda_1 = \lambda_2 = 1$. In this case, by projecting the ridge between WF_1 and WF_2 onto the plane of f_1 and f_2 we obtain the bisector of f_1 and f_2 , that is, the medial axis. However, if we choose $\lambda_1 = 1$, but we let $\lambda_2 > 1$ or $\lambda_2 < 1$, the projection of the ridge will move closer to or farther away from $f_2 = 0$ as illustrated in Figure 7b. Consequently, the reals λ_i act as weights and therefore the resulting R-function expression for the domain generates a weighted distance function (denoted here by WF). In fact,

$$\theta_i = \tan^{-1}(\lambda_i),$$

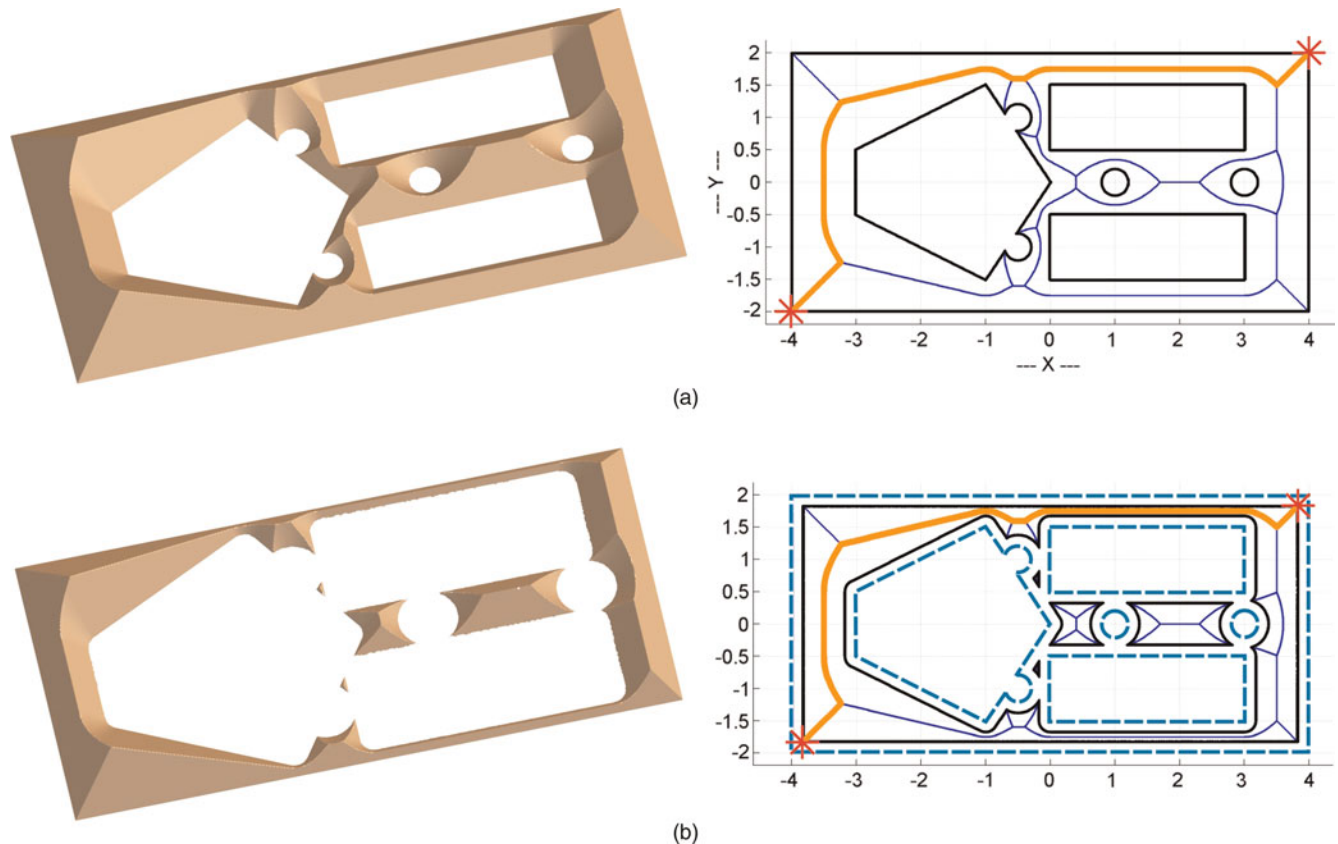


Fig. 9. (a) The trajectory planning of a point is shown. By taking different level sets of the implicit function induced by the R-function expression, we can plan the path of an object approximated by the (b) smallest enclosing ball. [A color version of this figure can be viewed online at journals.cambridge.org/aie]




Fig. 10. An environment with both static and dynamic obstacles moving independently. The planned path adapts to the changing environment. [A color version of this figure can be viewed online at journals.cambridge.org/aie]




Fig. 11. An environment with obstacles similar to the one in Figure 10b. The weight of the outer boundary is $\lambda = 0.2679$, which is equivalent to an angle of 15 degrees. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

where θ_i is the angle formed by $WF_i = 0$ and the plane (x, y) at the footpoint P_{f_i} on $f_i = 0$ as shown in Figure 7c. The fact that each scalar λ_i is constrained to be positive definite implies that the angle $\theta_i \in [0, \pi/2]$.

In other words, scalars λ_i act as degrees of freedom controlling the local attraction or repulsion of the skeleton by f_i . Furthermore, because each $\theta_i \in [0, \pi/2]$, the branch of the skeleton defined by two half-spaces $f_i \geq 0$ and $f_j \geq 0$ will lie in the region of the plane where both $f_i > 0$ and $f_j > 0$, that is, between $f_i > 0$ and $f_j > 0$.

In this paper, we assign a separate scalar λ to the distance function that defines each obstacle. In other words, each λ will act as a global weight for the corresponding obstacle as shown in Figure 7d. If the function defining each obstacle of the domain, which is represented by one single R-function expression, is weighted by a single weight λ , we conjecture that the resulting skeleton of the domain is homeomorphic to the medial axis, which in turn preserves the homotopy of the domain (Lieutier, 2003). In this case, the class of skeletons controlled by θ_i will have the same homotopy as that of the domain.

3. MEDIAL ZONES AS THICK SKELETONS

We have been using the principal system of R-functions given in Eq. (1) with a value of $\alpha = 1$. This, in turn, generates piecewise smooth distance functions as shown in Figure 3c, and the skeletons are obtained as the nondifferentiable points of the distance function as discussed in Sections 2.5 and 2.6. However, one important feature of R-functions is that they can generate continuous functions over the domain that are differentiable almost everywhere (for a detailed discussion, see Shapiro & Vossler, 1991; Shapiro, 2007). In fact, this is illustrated in Figure 3 for the principal system of R-functions of Eq. (1), which produces such differentiable functions by simply choosing a value of $\alpha < 1$.

For a given R-function expression corresponding to a given domain we define the *medial zones* as those points of \mathbb{E}^d that belong to the surface constructed with R-functions and are

“near” the crest of the surface,⁵ as illustrated in Figure 8. The “nearness” can be defined in several different ways and can be computed via Laplacian, Gaussian curvature, or other differential properties of the resulting surface. Note that the differential properties of the resulting surfaces can be controlled by the specific system of R-functions being used. Consequently, by adjusting the value of α and the “nearness” metric, one can “grow” or “shrink” these medial zones. The effect of changing α on the resulting medial zone is illustrated in Figure 8b and c, as detailed in Eftekharian and Ilieş (2010).

The implications to motion planning are important because these medial zones provide significantly increased flexibility of the resulting path compared to the case in which the paths are planned along the medial axis or skeleton of the same domain. Moreover, these medial zones can be computed with the same computational algorithms as those used to compute the skeletons described above. From a practical point of view, the paths planned based on the medial zones can be shorter and smoother than those planned based on skeletons, which is evident in the examples discussed in Section 4.

4. EXAMPLES

We illustrate the capabilities of our approach with several examples that contain both linear and nonlinear half-spaces in both 2-D and 3-D, as well as moving obstacles that merge with either other obstacles or the boundary of the domain $\partial\Omega$ to produce drastic topological changes of the environment. We implemented a commonly used finite-difference approximation of the Laplacian (Bovik, 2005) to extract the ridges and ravines of the weighted distance function.

Figure 9a shows the distance function of a domain and the corresponding shortest trajectory of a point between an initial and a final position. By taking a different level set of the distance function, we can plan within the same formulation the

⁵ Recall that the R-function expression generates the surface in the \mathbb{E}^{d+1} space.

path of an object approximated by the smallest enclosing ball as shown in Figure 9b.

The example in Figure 10 shows three instances during the evolution of an environment in which two obstacles are moving independently while merging and splitting with other obstacles. The planned path adapts to the changing environment as shown in the third row of images in the same figure.

The environment from Figure 10b is illustrated in Figure 11, but with a different path planning goal, namely, to have “sufficient” clearance to the boundary of Ω . The weight λ corresponding to the R-function expression defining the outer boundary of the domain is set to 0.2679 (equivalent to an angle

of 15 degrees; see Section 2.6). Note the difference between the resulting shortest path in Figure 11 and that of Figure 10b.

Figure 12 shows the path planning for a point within a domain with curved boundaries (Fig. 12a) and a second path of an object enclosed by a minimal ball while moving in the same environment. Note that the path planned for the object is different than the path of the point even though the space is the same. This is due to the fact that planning the path of the object is transformed into the path planning of a point by taking a different level set of the R-function expression describing the original environment. In this example the radius of the minimal enclosing ball is 0.3 m.

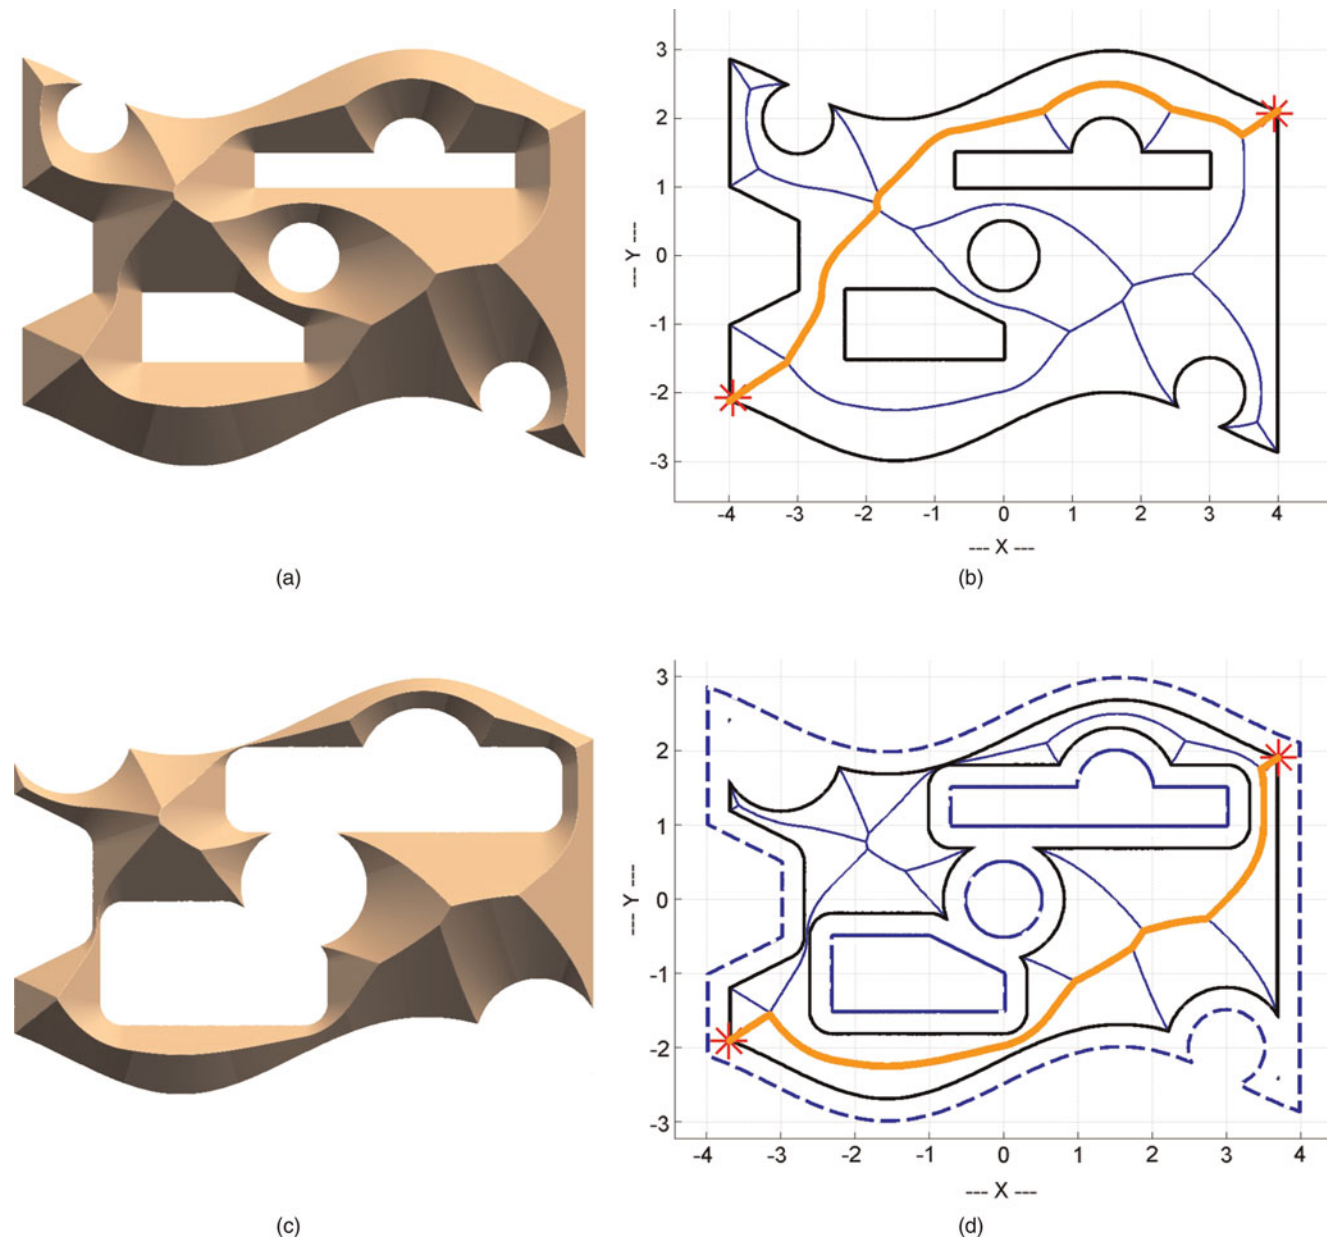


Fig. 12. (a) An environment with curved boundaries and the corresponding trajectory of a point and (b) the problem of planning a path of a moving object is transformed into the path planning of a point by taking a different level set of the distance function corresponding to the original environment. Note the difference in the two planned paths. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

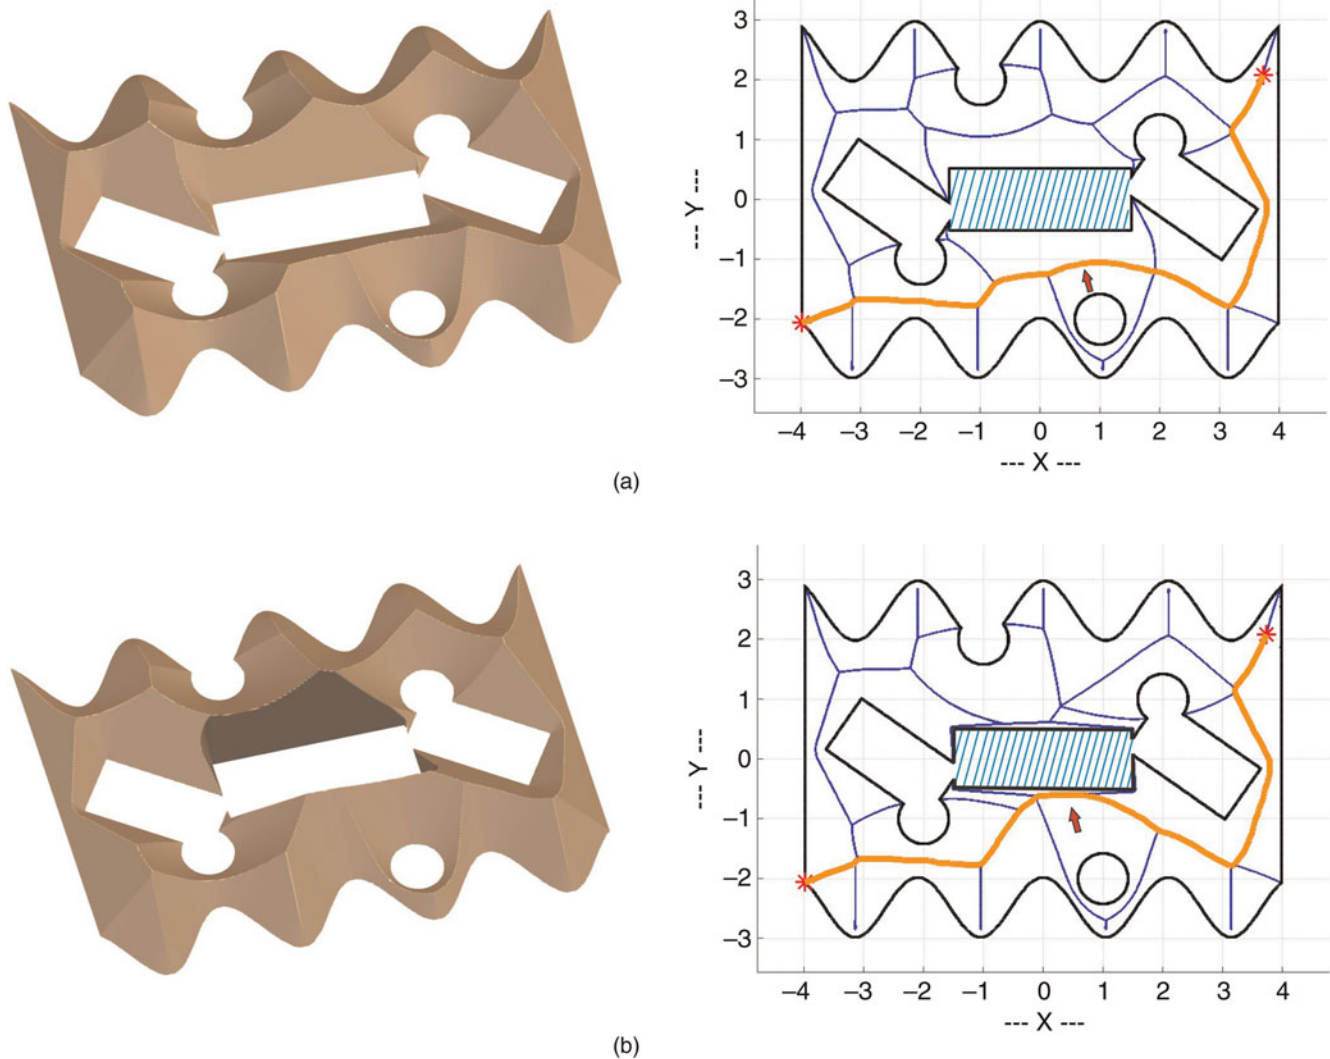


Fig. 13. Shortest path of a point within a domain with curved boundaries and five obstacles that was computed based on (a) the medial axis and (b) a skeleton obtained by weighing the distance function defining the rectangular (crosshatched) obstacle. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

A modified skeleton is shown in Figure 13 along with the computed path of a point within a domain with curved boundaries. The resulting roadmap computed based on the medial axis is shown in Figure 13a. Furthermore, the R-function expression defining the crosshatched rectangular obstacle has been weighted by a factor λ that corresponds to an angle of 85 degrees as discussed in Section 2.6. This effectively brings the skeleton, and therefore the computed path, closer to that particular obstacle. Note that there are five obstacles in this example, some of which are merged together, and that only one of the five corresponding distance functions is weighted. Clearly, other distance functions could be weighted as well depending on the specific path planning task being investigated with important applications in assembly, robotic, or navigation applications (as shown schematically in Fig. 1).

Figure 14 shows a domain containing multiple obstacles with free-form planar geometry. The three different cases illustrated

there correspond to three values of α , namely, $\alpha = 1$, $\alpha = 0.95$, and $\alpha = 0.8$, which result in the illustrated distance functions. For the case when $\alpha = 1$ illustrated Figure 14a the resulting skeleton is the medial axis, and the resulting shortest path is shown in orange (lighter color). By decreasing the value of α we increase the smoothness of the distance function, which in turn increases the “thickness” of the medial zones, which is illustrated in Figures 14b and c. Note that as α increases the length of the roadmap decreases, while its smoothness increases.

The R-function expression describing the domain of Figure 10b is illustrated in Figure 15a with a value of $\alpha = 0.90$, which results in the medial zone shown in Figure 15b. Observe that this new distance function effectively results in a smoother shortest path (Fig. 15b) that takes a different route than the shortest path shown in Figure 1b for the same domain.

Finally, Figure 16 shows several 3-D examples and the corresponding shortest paths. The 3-D domains shown in




Fig. 14. Shortest paths of a domain containing multiple obstacles with free-form planar geometry. The paths have been computed based on medial zones for three values of α . [A color version of this figure can be viewed online at journals.cambridge.org/aie]

Figure 16a and b contain one intermediate site that must be part of the path that is being planned. Note that a path along the medial axis that would satisfy these constraints does not exist. These examples contain two paths that do not satisfy the intermediate site constraint, which are planned inside two medial zones corresponding to values of $\alpha = 1$ and $\alpha = 0.9$, respectively. The same examples show a third

path, which, in each case, is the shortest path connecting the beginning and end configurations, passing through the intermediate site, while remaining inside a medial zone corresponding to a value of $\alpha = 0.9$. To achieve this, the distance function corresponding to the proximate obstacle (A) was modified in both examples as described in Section 2.6 and obstacle A “attracts” the resulting path. The example shown




Fig. 15. Medial zone and shortest path for the domain of Figure 1b for a value of $\alpha = 0.9$. The path computed in (b) is shown with dotted line for comparison purposes. [A color version of this figure can be viewed online at journals.cambridge.org/aie]




Fig. 16. Collision-free paths for several 3-D domains computed based on medial zones. The intermediate sites in (a) and (c) are shown as a points: the three paths correspond to values of $\alpha = 1$ and $\alpha = 0.9$; (b) the top view of the environment from (a); and (d) an environment with narrow passages. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

in Figures 16d illustrates another path planning scenario in an environment that contains many narrow passages and free-form obstacles.

5. CONCLUSION

This paper formulates a family of geometric skeletons that contains the medial axis as a special case, and explores their use in robotic path planning applications in highly dynamic and topologically evolving environments. We propose a strategy for constructing these skeletons that relies on constructive representations of shapes with R-functions that operate on real-valued half-spaces as logic operations. The flexibility provided by the underlying Boolean nature of the proposed framework makes this approach well suited to problems in-

volving substantial geometric and topologic changes of the environment. Note that our approach

1. is applicable for general closed, bounded, regular, and semianalytic domains;
2. results in a family of skeletons that contains the medial axis as a special case. The computed skeleton can be modified so that the computed path is attracted or repulsed by prescribed obstacles within the domain;
3. can compute the medial zones that intuitively represent thick skeletons of the domain. One important advantage of these zones in the context of motion planning applications is that they generate shorter and smoother paths than medial-axis path planning strategies for the same environment;

4. can handle problems with evolving environments or in which the environment is not fully known a priori because obstacles can be introduced/removed at any time during the planning stage and within the same problem formulation;
5. intrinsically supports local and parallel skeleton computations for domains with rigid or evolving boundaries due to the explicit mapping between the branches of the medial axis and the half-spaces bounding the environment;
6. can handle trajectory planning of points as well as path planning for 2-D/3-D objects within the same formulation due to the close relationship between R-function expressions and level sets. The simplest approach to extend the trajectory planning to solid objects is to approximate the moving object by its smallest enclosing ball, and take different level sets of the distance function defining the environment. This effectively offsets the boundary of the environment so that the outer boundaries of the environment shrink while, at the same time, the obstacles grow by the same radius of the disk enclosing the object; and
7. supports multiple representations of the input geometry;

We argue in this paper that the degrees of freedom that we used to control the subset of the domain in which we are looking for the shortest path afford significant flexibility in computing paths that satisfies specific geometric constraints that are common in geometric reasoning and motion planning applications. In addition, the attractive and repulsive factors λ_i provide a simple and elegant mechanism to deform the skeleton so that the resulting path gets closer to or farther away from specific sites or obstacles in the environment.

We conjectured that the family of skeletons defined by weights λ as defined in Section 2.4 is homeomorphic to the medial axis, which implies that the skeletons in the family are homotopic to the domain. This property coupled with the fact that the family of skeletons is really defined by pseudodistance functions is crucial in all potential applications such as unmanned vehicles navigation systems, robotics, autonomous assembly systems, and shape representation and recognition.

ACKNOWLEDGMENTS

This work was supported in part by National Science Foundation Grant CMMI-0555937 and CAREER Awards CMMI-0644769, CMMI-0856401, and CNS-0927105.

REFERENCES

- Amenta, N., Choi, S., & Kolluri, R.K. (2001). The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications* 19, 127–153.
- Attali, D., Boissonnat, J.-D., & Edelsbrunner, H. (2008). Stability and computation of medial axes—a state-of-the-art report. In *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration, Mathematics and Visualization* (Möller, T., Hamann, B., & Russel, R., Eds.). Berlin: Springer-Verlag.
- Blum, H. (1967). A transformation for extracting new descriptions of shape. *Proc. Models for the Perception of Speech and Visual Form*, pp. 362–380.
- Blum, H., & Nagel, R. (1978). Shape description using weighted symmetric axis features. *Pattern Recognition* 10, 167–180.
- Bookstein, F. (1979). The line-skeleton. *Computer Graphics and Image Processing* 11(2), 123–137.
- Bovik, A.C. (2005). *Handbook of Image and Video Processing (Communications, Networking and Multimedia)*. Orlando, FL: Academic Press.
- Brandt, J.W., & Algazi, V.R. (1992). Continuous skeleton computation by Voronoi diagram. *CVGIP: Image Understanding* 55(3), 329–338.
- Buchele, S.F., & Crawford, R.H. (2003). Three-dimensional halfspace constructive solid geometry tree construction from implicit boundary representations. *Proc. 8th ACM Symp. Solid Modeling and Applications*, pp. 135–144. New York: ACM Press.
- Canny, J. (1993). Computing roadmaps of general semi-algebraic sets. *Computer Journal* 36(5), 504–514.
- Cao, L., & Liu, J. (2008). Computation of medial axis and offset curves of curved boundaries in planar domain. *Computer-Aided Design* 40(4), 465–475.
- Chazal, F., & Soufflet, R. (2004). Stability and finiteness properties of medial axis and skeleton. *Journal of Dynamical and Control Systems* 10(2), 149–170.
- Chen, J., Shapiro, V., Suresh, K., & Tsukanov, I. (2007). Shape optimization with topological changes and parametric control. *International Journal for Numerical Methods in Engineering* 71, 313–346.
- Choi, H.I., Choi, S.W., & Moon, H.P. (1997). Mathematical theory of medial axis transform. *Pacific Journal of Mathematics* 1(181), 57–88.
- Choset, H. (2000). Coverage of known spaces: the boustrophedon cellular decomposition. *Autonomous Robots* 9, 247–253.
- Culver, T., Keyser, J., & Manocha, D. (2004). Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design* 21(1), 65–98.
- Damon, J. (2005). Determining the geometry of boundaries of objects from medial data. *International Journal of Computer Vision* 63, 45–64.
- Dasgupta, S., Papadimitriou, C.H., & Vazirani, U. (2008). *Algorithms*. New York: McGraw-Hill.
- Diaz de León, S.J.L., & Sossa, A.J.H. (1998). Automatic path planning for a mobile robot among obstacles of arbitrary shape. *IEEE Transactions on Systems, Man, and Cybernetics, Part B* 28(3), 467–472.
- Dey, T.K., Woo, H., & Zhao, W. (2003). Approximate medial axis for CAD models. *Proc. 8th ACM Symp. Solid Modeling and Applications: ISM '03*, pp. 280–285. New York: ACM.
- Dijkstra, E.W. (1959). A note on two problems in connection with graphs. *Numerische Mathematik* 1, 269–271.
- Dobkin, D., Guibas, L., Hershberger, J., & Snoeyink, J. (1988). An efficient algorithm for finding the CSG representation of a simple polygon. *Proc. 15th Annual Conf. Computer Graphics and Interactive Techniques: SIGGRAPH '88*, pp. 31–40. New York: ACM.
- Du, H., & Qin, H. (2004). Medial axis extraction and shape manipulation of solid objects using parabolic PDEs. *Proc. 9th ACM Symp. Solid Modeling and Applications: SM '04*, pp. 25–35. Aire-la-Ville, Switzerland: Eurographics Association.
- Eftekharian, A., & Ilieş, H. (2009). Distance functions and skeletal representations of rigid and non-rigid planar shapes. *Computer-Aided Design* 41(12), 865–876.
- Eftekharian, A., & Ilieş, H. (2010). *Medial zones: formulation, properties and applications*. Technical Report, University of Connecticut.
- Foskey, M., Garber, M., Lin, M.C., & Manocha, D. (2001). A Voronoi-based hybrid motion planner. *Proc. IEEE/RSJ Int. Conf. Intelligent Robots & Systems*, pp. 55–60.
- Fraichard, T. (1993). Dynamic trajectory planning with dynamic constraints: a “state-time space” approach. *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems '93: IROS '93*, Vol. 2, pp. 1393–1400.
- Garber, M., & Lin, M.C. (2002). Constraint-based motion planning for virtual prototyping. *Proc. 7th ACM Symp. Solid Modeling and Applications: SMA '02*, pp. 257–264. New York: ACM.
- Geraerts, R., & Overmars, M.H. (August 2007). Creating high-quality paths for motion planning. *International Journal of Robotics Research* 26(8), 845–863.
- Goh, W.-B. (2008). Strategies for shape matching using skeletons. *Computer Vision and Image Understanding* 110, 326–345.
- Gomes, J., & Faugeras, O. (2000). Reconciling distance functions and level sets. *Journal of Visual Communication and Image Representation* 11(2), 209–223.

- Hoff, K., III, Culver, T., Keyser, J., Lin, M., & Manocha, D. (2000). Interactive motion planning using hardware-accelerated computation of generalized Voronoi diagrams. *Proc. IEEE Int. Conf. Robotics and Automation, 2000: ICRA '00*, Vol. 3.
- Ishikawa, S. (1991). A method of indoor mobile robot navigation by using fuzzy control. *Proc. IEEE/RSJ Int. Workshop on Intelligent Robots and Systems: IROS '91*, Vol. 2, pp. 1013–1018.
- Jiang, K., Seneviratne, L.D., & Earles, S.W.E. (1999). A shortest path based path planning algorithm for nonholonomic mobile robots. *Journal of Intelligent and Robotic Systems 24(4)*, 347–366.
- Kallman, M., & Mataric, M. (2004). Motion planning using dynamic roadmaps. *Proc. 2004 IEEE Int. Conf. Robotics and Automation: ICRA '04*, Vol. 5, pp. 4399–4404.
- Kavraki, L.E., Vestka, P., Latombe, J.-C., & Overmars, M.H. (1996). Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation 12*, 566–580.
- Kimia, B.B., & Tamrakar, A. (2002). The role of propagation and medial geometry in human vision. *Proc. 2nd Int. Workshop on Biologically Motivated Computer Vision: BMCV '02*, pp. 219–229. London: Springer-Verlag.
- Kimmel, R., Shaked, D., Kiriyati, N., & Bruckstein, A. (1995). Skeletonization via distance maps and level sets. *Computer Vision and Image Understanding 62(3)*, 382–391.
- Latombe, J.-C. (1991). *Robot Motion Planning*. Norwell, MA: Kluwer Academic.
- LaValle, S.M. (2006). *Planning Algorithms*. Cambridge: Cambridge University Press.
- Lee, J.Y., & Choset, H. (May 2005). Sensor-based planning for a rod-shaped robot in three dimensions. *International Journal of Robotics Research 24*, 343–383.
- Lieutier, A. (2003). Any open bounded subset of \mathbb{R}^n has the same homotopy type than its medial axis. In *Proc. 8th ACM Symp. Solid Modeling and Applications: SM '03*, pp. 65–75. New York: ACM.
- Lingelbach, F. (2004). Path planning using probabilistic cell decomposition. *Proc. IEEE Int. Conf. Robotics and Automation*.
- Liu, T.-L., Geiger, D., & Kohn, R.V. (1998). Representation and self-similarity of shapes. *Proc. 6th Int. Conf. Computer Vision: ICCV '98*, pp. 1129. Washington, DC: IEEE Computer Society.
- Luo, Z., Tong, L., Wang, M.Y., & Wang, S. (2007). Shape and topology optimization of compliant mechanisms using a parameterization level set method. *Journal of Computer Physics 227*, 680–705.
- Masehian, E., & Katebi, Y. (2007). Robot motion planning in dynamic environments with moving obstacles and target. *International Journal of Mechanical Systems Science and Engineering 1(1)*, 20–25.
- Neus, M., & Maouche, S. (2005). Motion planning using the modified visibility graph. *Proc. IEEE Int. Conf. Systems Man and Cybernetics*, Vol. 4, pp. 651–655.
- Pizer, S.M., Siddiqi, K., Székely, G., Damon, J.N., & Zucker, S.W. (2003). Multiscale medial loci and their properties. *International Journal of Computer Vision 55(2–3)*, 155–179.
- Quadros, W., Shimada, K., & Owen, S. (2004). Skeleton-based computational method for the generation of a 3D finite element mesh sizing function. *Engineering with Computers 20(3)*, 249–264.
- Ramamurthy, R., & Farouki, R.T. (1999). Voronoi diagram and medial axis algorithm for planar domains with curved boundaries—II: detailed algorithm description. *Journal of Computers and Applied Mathematics 102(2)*, 253–277.
- Requicha, A. (1980). Representations for rigid solids: theory, methods and systems. *Computing Surveys 12(4)*, 437–463.
- Rimon, E., & Koditschek, D. (1992). Exact robot navigation using artificial potential functions. *IEEE Transactions on Robotics and Automation 8(5)*, 501–518.
- Sacks, E., Pisula, C., & Joskowicz, L. (1999). Visualizing 3d configuration spaces for mechanical design. *IEEE Computer Graphics and Applications 19(5)*, 50–53.
- Sampl, P. (2000). Semi-structured mesh generation based on medial axis. *Proc. IMR*, pp. 21–32.
- Sampl, P. (2001). Medial axis construction in three dimensions and its application to mesh generation. *Engineering with Computers 17*, 234–248.
- Sebastian, T., Klein, P., & Kimia, B. (2004). Recognition of shapes by editing their shock graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence 26(5)*, 550–571.
- Shah, J. (2005). Gray skeletons and segmentation of shapes. *Computer Vision and Image Understanding 99*, 96–109.
- Shaham, A., Shamir, A., & Cohen-Or, D. (2004). Medial axis based solid representation. *Proc. 9th ACM Symp. Solid Modeling and Applications*.
- Shapiro, V. (1991). *Theory of R-Functions and Applications: A Primer*, Technical Report TR91-1219. Ithaca, NY: Cornell University, Computer Science Department.
- Shapiro, V. (2001). A convex deficiency tree algorithm for curved polygons. *International Journal of Computational Geometry and Applications 11(2)*, 215–238.
- Shapiro, V. (2007). Semi-analytic geometry with R-functions. *Acta Numerica 18*, 239–303.
- Shapiro, V., & Vossler, D.L. (1991). Construction and optimization of CSG representations. *Computer-Aided Design 23(1)*, 4–20.
- Shapiro, V., & Vossler, D.L. (1993). Separation for boundary to CSG conversion. *ACM Transactions on Graphics 12(1)*, 35–55.
- Sherbrooke, E., Patrikalakis, N., & Brisson, E. (1995). Computation of the medial axis transform of 3-D polyhedra. *Proc. 3rd ACM Symp. Solid Modeling and Applications*, pp. 187–200. New York: ACM.
- Siddiqi, K., Bouix, S., Tannenbaum, A., & Zucker, S.W. (2002). Hamilton–Jacobi skeletons. *International Journal of Computer Vision 48(3)*, 15–231.
- Siddiqi, K., Shokoufandeh, A., Dickinson, S.J., & Zucker, S.W. (1999). Shock graphs and shape matching. *International Journal of Computer Vision 35(1)*, 13–32.
- Snaveley, N., Garg, R., Seitz, S., & Szeliski, R. (2008). Finding paths through the world's photos. *ACM Transactions on Graphics 27(3)*, 15.
- Suresh, K. (2003). Automating the CAD/CAE dimensional reduction process. *Proc. 8th ACM Symp. Solid Modeling and Applications: SM '03*, pp. 76–85. New York: ACM.
- Takaoka, T. (2005). An $O(n^3 \log \log n / \log n)$ time algorithm for the all-pairs shortest path problem. *Information Processing Letters 96(5)*, 155–161.
- Tor, S.B., & Middleditch, A.E. (1984). Convex decomposition of simple polygons. *ACM Transactions on Graphics (TOG) 3(4)*, 244–265.
- van den Berg, J., & Overmars, M. (2005). Roadmap-based motion planning in dynamic environments. *IEEE Transactions on Robotics 21(5)*, 885–897.
- van Eede, M., Macrini, D., Telea, A., Sminchisescu, C., & Dickinson, S. (2006). Canonical skeletons for shape matching. *Proc. 18th Int. Conf. Pattern Recognition ICPR 2006*, Vol. 2, pp. 64–69.
- Wilmarth, S.A., Amato, N.M., & Stiller, P.F. (1999). Motion planning for a rigid body using random networks on the medial axis of the free space. *Proc. 15th Annual Symp. Computational Geometry: SCG '99*, pp. 173–180. New York: ACM.
- Xidias, E.K., Azariadis, P.N., & Aspragathos, N.A. (2007). Two-dimensional motion-planning for non-holonomic robots using bump-surface concept. *Computing 79*, 109–118.

Ata Eftekharian is a Postdoctoral Fellow in the Automated Design Lab, Department of Mechanical Engineering, at the University of Texas at Austin. Dr. Eftekharian has MS and PhD degrees in mechanical engineering from Sharif University of Technology and University of Connecticut, respectively. His research interests include geometric modeling, computer-aided design and manufacturing systems, and artificial intelligence. His research work mainly focuses on developing new theories and techniques for automating synthesis of mechanical artifacts, path planning, and topics related to reasoning about manufacturability and design changes.

Horea Ilies is an Associate Professor of mechanical engineering and computer sciences at the University of Connecticut, where he established the Computational Design Laboratory. Prior to that, he spent several years with Ford Motor Company in research, manufacturing, and product design and development. His current research interests include geometric and physical computing, shape synthesis and geometric reasoning, as well as biomolecular engineering. Horea received the NSF CAREER award in 2007.