

## Automatically Generating Virtual Guided Tours

Tsai-Yen Li, Jyh-Ming Lien, Shih-Yen Chiu, and Tzong-Hann Yu  
Computer Science Department  
National Chengchi University,  
Taipei, Taiwan, R.O.C.  
{li, s8415, s8433, s8410}@cs.nccu.edu.tw

### Abstract

*After the introduction of VRML, 3D web browsing has become a popular form of networked virtual reality. However, it is still a great challenge for a novice user equipped with a regular desktop PC to navigate in most virtual worlds of moderate complexity. We think the main problem is due to the fact that a user usually uses a 2D mouse to provide low-level navigation control but the display frame rate is not high enough for this servo loop. In this paper, we consider an alternative metaphor of allowing a user to specify locations of interests on a 2D-layout map and let the system automatically generate the animation of guided tours in virtual architectural environments. Specifically, we aim to generate animations of customizable tour paths and its associated human/camera motions in an on-line manner according to high-level user inputs. We describe an auto-navigation system, in which several efficient path-planning algorithms adapted from robotics are used. This system has been implemented in Java and adopts common VRML browsers as its 3D interface. We also use the geometric model of our departmental building as an example to demonstrate the efficiency and effectiveness of the system.*

### 1. Introduction

Virtual Reality has become a popular form of publishing 3D data on the network. It is partly due to the ever-growing computing power and 3D graphics acceleration hardware available for personal computers. In addition, the wide acceptance on the standardization of 3D data formats, such as Virtual Reality Modeling Language (VRML)[23], also speeds up the development of networked VR. However, there still exist problems preventing this form of networked VR from being effectively used on a regular desktop PC. For example, manipulating a complex scene with a 2D mouse is not an easy task for a novice user, and a high frame rate is required for the user to get familiar with the metaphor. A high frame rate can be achieved by using a good 3D graphics acceleration card and a carefully designed scene. Nevertheless, as the com-

puting power on PC's increases, a user would usually expect to be able to navigate in a more realistic (and more complex) virtual world. Therefore, we believe that the computing power will never catch up with user demands on high-quality graphics. Consequently, in many applications, such as interactive virtual tours, a different metaphor is needed for more effective navigation.

In this paper we attempt to use machine intelligence as a means to increase the level of control that a user need to provide for navigating in a virtual world effectively. For example, when a user visits a site in the real world with full services, he/she may only need to express his/her interests by specifying locations on a 2D-layout map. A tour guide should then plan an optimal route tailored to this visitor and take him/her through these locations of interests. During the customized tour, instead of following the guide's trace, the visitor may follow the tour guide in a more leisurely way such that keeping the visibility with the tour guide is the only requirement.

We will describe an implemented networked VR system aiming to provide the tour-guiding service envisioned in the previous paragraph. We examine three main modules required to provide the services: *customizable tour planning*, *automated humanoid animation*, and *intelligent camera positioning*. A user of the system only needs to specify locations of interest on a 2D-layout map and select a traversing mode, the system will generate a good (if not optimal) tour path passing these locations. The animation of a human tour guide will also be generated automatically in real time according to the tour path. In addition, the system will plan the motion of the camera representing the user's viewpoint such that the tour guide is always kept inside the view model of the camera.

In the next section, we will present the problems involved in providing the above services and related work in the literature for solving these problems. We give a more formal description of these problems in Section 3 and propose our approaches in Section 4 to solve them in an on-line manner. System architecture and implementation details will then be presented in Section 5. Experimental results and illustrative examples will be shown in Section 6. Conclusions and future extensions of this work are then

given in Section 7.

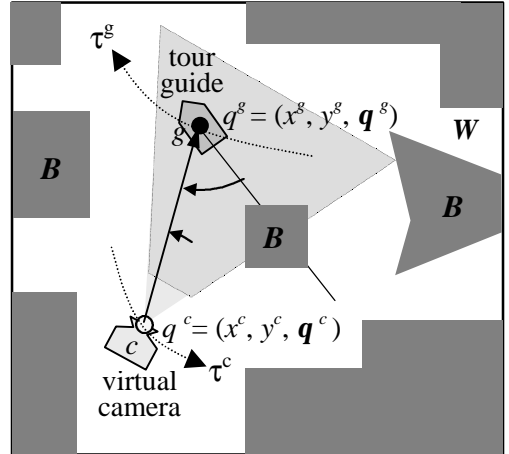
## 2. Related Work

There have been many researches in intelligent agents that try to provide various forms of guidance to a user in a virtual environment[4][18]. However, we have not seen systems that incorporate geometric reasoning, such as path planning or camera tracking, into the application of virtual guided tours. In the scenario described in the previous section, we need to consider three problems that are not accounted for in the previous systems. First, the path followed by the tour guide should be planned in an on-line fashion with a given user specification. Second, the human walking animation for the tour guide should be generated in real time according to the tour path. Third, the tracking path of the virtual camera representing the user should also be planned according to the tour path. The first and the third problems are path-planning problems while the second problem belongs to the category of simulation.

The first problem is to generate a good (preferably optimal) tour path that takes the tour guide through all desired locations without colliding with obstacles in the environment[12]. The problem of planning a collision-free path between two configurations is the so-called ‘‘Piano Mover’s Problem’’ in the robotics literature, and it has been widely studied in the past two decades[9]. Researches have shown that the general path-planning problem is PSPACE-hard, and its complexity grows exponentially in the number of degrees of freedom (DOF) that the moving object has[17]. When there are several locations to traverse, the problem of determining an optimal traversing sequence becomes the classical ‘‘Travelling Salesperson’s Problem (TSP)’’, which is known to be NP-complete[3].

The second problem is about generating realistic animation of human walking gaits with various step lengths. There have been many researches in the graphics literature about generating humanoid animation. Some of them use pure kinematics models while others include dynamics considerations as well[6][20][19]. Recent researches also focus on reusing captured data on different environmental settings through wrapping or transforming in a different domain[22][24]. In our scenario, the tour guide walks with regular steps of various lengths. Thus, no special walking gaits need to be considered. The only requirement is that the generation of simulation needs to be very efficient in order to be used in a real-time environment.

The third problem is about moving the camera in such a way that the tour guide is always kept inside its view cone. That is, the generated camera motion not only needs to be collision-free with the obstacles in the environment, the view between the camera and the guide should also be clear all the time. The problem is similar to the Intelligent



**Figure 1: Cartesian-space and view-space configurations of the tour guide and the camera.**

Observer problem considered in robotics[2][7][11]. In the system described in [11], an optimal tracking path can be computed with a dynamic programming approach. However, for a coarse resolution of workspace, it takes about 20 seconds to generate such a path on a modern workstation. This planning efficiency is not good enough for our on-line application.

## 3. Problem Description

Since computational efficiency is the most important concern in our interactive auto-touring application, the path-planning problems need to be simplified. For example, since the guided tours considered in our system are all in architectural environments, it is reasonable to assume that objects in the scene can be projected onto a 2D workspace, as shown in Figure 1, where the planning problem is defined.

### 3.1. The path-planning problems

We assume that both the tour guide and the virtual camera have three degrees of freedom (3-DOF) on a plane, and their geometry can be reasonably represented by 2D polygonal objects. These two objects move in a bounded 2D workspace,  $W$ , cluttered with polygonal obstacles,  $B$ . Let the configurations of the tour guide  $g$  and the camera  $c$  be denoted by  $q^g = (x^g, y^g, \mathbf{q}^g)$  and  $q^c = (x^c, y^c, \mathbf{q}^c)$ , respectively. The *configuration spaces (C-spaces)* for these two objects are composed of all possible configurations of  $q^g$  and  $q^c$ , respectively[14]. The composite space of these two C-spaces is denoted by  $X$ , where each composite configuration is denoted by  $q = (x^g, y^g, \mathbf{q}^g, x^c, y^c, \mathbf{q}^c)$ . The objective of the auto-touring system is to generate a legal

path  $\mathbf{t}$  in  $X$  connecting an initial configuration  $q_i$  and a final configuration  $q_f$ . A configuration is *legal* if and only if  $g$  and  $c$  do not overlap with  $B$  or with each other, and  $g$  is inside the view cone of  $c$  at any time as shown in Figure 1.


The solution to our tour-planning problem resides in this composite C-space defined above. However, it is not likely that we can design an on-line planner that can solve the problem in a space of such high dimensionality. (Most on-line planners exist for problems with less than 4-DOF[13].) Fortunately, in our tour-planning problem, the motions of the tour guide and the camera are not without dependency. The path for the tour guide should be determined before the camera path. Therefore, a reasonable approach would be to plan a tour path  $\mathbf{t}^g$  first and then generate the camera path  $\mathbf{t}^c$  according to  $\mathbf{t}^g$ .

After  $\mathbf{t}^g$  is determined, we can plan the camera motion by connecting them with an imaginary line segment dragging the camera as the tour guide moves. Therefore, instead of planning in the camera's Cartesian space  $(x, y, \mathbf{q})$ , we re-parameterize the camera configuration with an equivalent representation,  $(l, \mathbf{f}, \mathbf{j})$ , relative to the tour guide, as shown in Figure 1. These parameters, defined simply for planning convenience, are called *view distance*, *tracking angle* and *view angle*, respective. Since  $\mathbf{t}^g$  is a function of time, the planning space for  $\mathbf{t}^c$  now becomes a 4D *configuration-time space (CT-space)* defined by  $(t, l, \mathbf{f}, \mathbf{j})$ .

### 3.2. The optimal sequence problem

One of the challenges in this auto-touring system is on generating an optimal traversing sequence through the user-specified locations of interests. An illustrative path traversing through all user-specified locations is shown in Figure 2. This problem belongs to the classical NP-complete TSP problem, where no efficient algorithms are found so far. Nevertheless, there exist approximate algorithms that can find near-optimal solutions in polynomial time if the Euclidean distances between locations can be determined[3]. However, for a workspace cluttered with obstacles, the length of a collision-free path is the real distance between two locations, and this information is not known prior to the search process. Therefore, it is difficult to find an optimal solution in general. However, the auto-touring problem under consideration may still contain some characteristics that allow an efficient implementation for sub-optimal solutions.

### 3.3. The humanoid animation problem



**Figure 2: a sample tour plan traversing all user-specified locations marked with x. (Obstacles are in grey.)**

Our system aims to generate realistic humanoid animation in a real-time fashion. This is because the tour path is generated according to on-line user specifications, the humanoid animation can not be prepared ahead of time. Instead, after a feasible tour path has been determined, the path is parameterized into a list of footprints with various step lengths and possibly at different elevations. Given these footprints, the system is to generate 3D human motions adapting to local obstacles in real time with appropriate joint angles for a given human kinematics model.

## 4. The Proposed Approaches

In this section, we propose our practical approaches to the computationally difficult problems described in the previous section. The objectives of these approaches is to bring the planning or simulation activities on-line for interactive applications. Simplifications are usually required to achieve this goal but we will show that through careful design and appropriate simplifications, the results could still be quite satisfactory in terms of efficiency and quality.

The planning approach we take in this paper is similar to the distributed approach described in [1]. The planning spaces are represented by uniform grids of certain resolution, and the planning algorithms are variations of the best-first search algorithm with different “best” criteria.

### 4.1. Decoupled approach for path planning

As mentioned in the previous section, although the overall planning space is a six-dimensional space, the dimensions are not of independent. Therefore, we take a decoupled approach by decomposing the problem into two smaller sub-problems of planning for two objects of 3-DOF each and solving them in sequence[5]. That is, we first plan the tour path in the tour guide's C-space (3D) and then plan the camera path in the camera's CT-space (4D) with the previously planned tour path as a constraint.

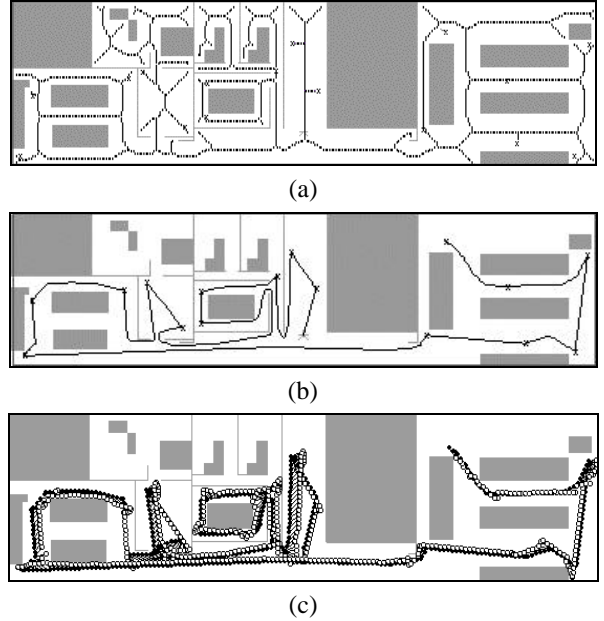
There exist efficient algorithms in robotics that can generate collision-free paths in fractions of a second for problems of less than 5D[1][13]. However, in the worst-case situations the running time could still be tens of seconds or even minutes. In addition, the length of the overall path in our application is typically much longer than the one considered in robotics. Therefore, further simplifications or decompositions are needed for better on-line performance. For each of the two sub-problems above, we further decompose their search space according to the moving object’s geometric characteristics.

For the tour-planning problem, we assume that the geometry of tour guide can be approximated by an enclosing circle. We can then reduce the dimensionality of the C-space from three to two due to the geometric symmetry. The C-space can be explicitly computed by growing the workspace obstacles with the radius of the enclosing circle[10]. This is a conservative but reasonable assumption because in our application a planned tour does not normally guide a visitor to pass through narrow passages. The tour paths generated in this reduced C-space consist of 2D configurations  $(x, y)$  only. We add the decoupled  $\mathbf{q}$  component of each configuration back into the path in a postprocessing step. In this step, we assume that a tour guide always faces the tangential direction of the path. At the locations where large discrepancies in  $\mathbf{q}$  occur, we add additional steps into the path to make the rotational transition smooth.

For the camera-planning problem, we also take a decoupled approach to reduce the dimensionality of the search space. For the CT-space  $(t, l, \mathbf{f}, \mathbf{j})$  defined in the previous section, we lock the view angle,  $\mathbf{j}$ , between the camera and the tour guide such that the tour guide is always at the center of the camera’s field of view. With this assumption, we then search for a collision-free path in the reduced CT-space connecting an initial configuration  $q_i=(t_i, l_i, \mathbf{f}_i)$  to any legal final configuration  $q_f=(t_f, *, *)$ , where \* represents any legal values. After a feasible tracking path has been found, we unlock and fill in the  $\mathbf{j}$  dimension in such a way that the camera rotations are minimized. More detail description of the tracking motion planning can be found in [14].

#### 4.2. Greedy approach for optimal sequencing


The difficulties of determining an optimal traversing sequence in our problem lie on the unknown distances between locations and the large number of possible traversing orders in general. We overcome these difficulties by taking a greedy approach, which traverses on a skeleton (discrete Voronoi diagram) built in the free space of the aforementioned 2D C-space. An example of the skeleton is shown in Figure 3(a). The skeleton captures



**Figure 3: an example of (a) freespace skeleton, (b) tour path, and (c) the corresponding camera path, where the guide and the camera are depicted as solid dots and hollow circles, respectively.**

the topological information of the freespace and can be computed automatically for a workspace bitmap in a preprocessing step[9]. All locations to be traversed are connected to the skeleton on the fly and treated as part of the skeleton before the planning starts. Unlike the classical path-planning problem where the initial and goal configurations are well defined prior to planning, we do not have a definite goal to reach in our problem. We start the search from the initial configuration and allow the search to propagate in a breadth-first fashion along the skeleton until it reaches a nearest goal location that has not been visited before. We repeat this process until all locations are visited. Since the search is limited on the skeleton, the search space is considerably reduced. Although the resulting tour path is not guaranteed to be optimal, this greedy approach works pretty well in the examples we have tested. Examples of the tour path and the camera path generated using these approaches are shown in Figures 3(b) and 3(c), respectively. In Figure 3(c), the tour guide (solid dot) and the virtual camera (hollow circle) are connected by a line segment, which remains collision free with the obstacles all the time along the tour. Detail data for planning will be reported in Subsection 6.2.

#### 4.3. Constrained kinematics approach for real-time humanoid simulation



**Figure 4: side views of the lower-body animations produced by the real-time humanoid simulator for different step sizes and at different elevations.**

In this subsection we describe how we generate the lower-body animation of a humanoid model. This walking simulation needs to be able to adapt to possible environmental constraints imposed on foot locations. We re-parameterize the path computed in the tour planner into a list of footsteps of various lengths possibly at different elevations. Due to the real-time requirement of our application, we do not consider dynamics models in our system. Instead, we use a constraint-based pure kinematics approach that unifies the walking models for normal, uphill, and downhill situations into one.

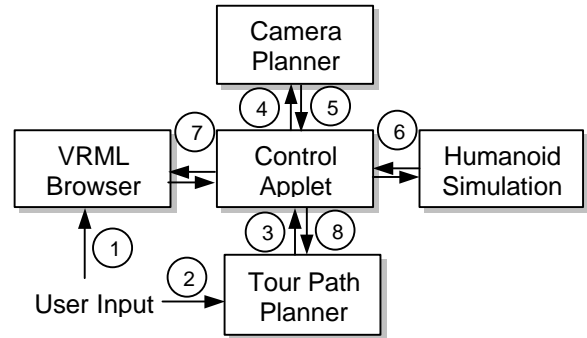
According to [6], a gait cycle consists of a transfer duration and a support duration. We further divide the cycle into four phases of different characteristics and identify the four key frames between them. In each phase, an appropriate dominant parameter either in the joint space or in the Cartesian space is chosen to drive the walking gait. These four key frames are identified as follows.

- (1) Both legs are straight down and one leg (the front leg) starts to kick off.
- (2) The front leg touches the ground.
- (3) The back leg starts to lift off.
- (4) The center of gravity is on top of the front leg again.

The cycle repeats with the front leg and the back leg switch alternatively. Between two key frames, the lower-body motions are interpolated according to the following two principles:

- I. For the leg on the ground, motions are interpolated in the joint space, and forward kinematics is used to compute the hip location.
- II. For the leg in the air, a collision-free cubical curve generated according to the local ground geometry is used for interpolation in Cartesian space. Inverse kinematics is used to compute the joint values.

In addition to its computational efficiency, one of the advantages with this kind of approach is on the ability to generate realistic motions adapting to different step lengths and foot elevations. Examples generated by this



**Figure 5: system architecture and data flow among different modules in the tour-guiding system**

simulator are shown in Figure 4(a) and 4(b).

## 5. System Architecture and Implementation

In this section, we describe how we implemented the planning or simulation modules and how these different modules work together to provide the auto-touring service.

### 5.1. System architecture

The system is comprised of five logical modules as depicted in Figure 5: VRML display module, control module, tour-planning module, camera-planning module, and humanoid simulation module. Only the VRML and tour-planning modules contain graphical user interface. A user specifies the locations of interests by pointing directly on a layout map presented by the tour-planning applet. The user can also specify preferences such as using default locations, choosing a preferred traversing order, etc. through this interface. The final guided tour is presented to the user via a common VRML browser in a regular web page.

A typical data flow in the system, as shown in Figure 5, is described as follows.

- (1) manually navigating in the VRML browser,
- (2) specifying locations of interests via the 2D interface and starting to generate the customized tour when ready,
- (3) passing the generated tour path to the control module,
- (4) passing the tour path to the camera planner for generating the camera motion,
- (5) returning the camera motion back to the control module,

- (6) passing the tour path to the humanoid simulation module to generate appropriate joint motions for walking simulation,
- (7) passing the motions of the tour guide and the camera to the VRML browser to create animation of the virtual tour,
- (8) updating the locations of the tour guide and camera in the 2D interface.

The animation stops when the system has finished traversing all desired locations.

## 5.2. Implementation

Except for the VRML display module, all modules described in the previous subsections are implemented as Java applets. These applets communicate with each other through the scripting model of the hosting web browser. The control applet communicates with the VRML browser via the External Authoring Interface (EAI) [16].


Although the 2D floor map used in the planners can be generated directly from the VRML models, it is currently read in from a data file created separately. The dimensions of the tour guide used in the walking gait simulator are computed according to the standard humanoid model in the VRML file. This model is obtained dynamically via the EAI module and can therefore be easily replaced with other avatars constructed in accordance with the humanoid specification[21].

## 6. Experimental Results

All of the running times reported in this paper were taken on a 450MHz Pentium II PC running MS Windows NT. In the current system, we have used our departmental building as an example to demonstrate the efficiency and usability of the system.

### 6.1. Graphical user interface

The system provides a graphical user interface in a web browser running on regular desktop PC's. A snapshot of the graphical user interface of the system is shown in Figure 6. The VRML browser is at the top while the 2D floor map is shown at the bottom. The snapshot was taken at our personal computer laboratory. The size of the VRML file is about 1.3MB and the frame rate is about 0.5 to 4 frames per second. On a regular PC, it is indeed a pain to navigate in a virtual scene of this complexity. However, this is exactly the main reason that motivates our auto-touring system in the first place. During the animation of the tour, a user can interact with the animation directly on the VRML browser or through the 2D interface. In particular, in the 2D interface (at the lower right



**Figure 6: a snapshot of the graphical user interface in the touring system**

corner of Figure 6) a user can adjust the head orientation horizontally or vertically independent of the body motion (which is always on a plane).

### 6.2. Planning efficiencies

The search spaces of all path planners in this system are represented as multidimensional bitmaps. The workspace bitmap resolution used in the tour planner is 560x150 while the maximal allowable rotation in a walking step is 10 degrees. In the camera planner the resolution for the last two parameters in the CT-space ( $t, l, f$ ) is 50x100, and the resolution in the  $t$  dimension depends on the number of steps in the path generated by the tour planner. In the example shown in Figure 3, there are 1112 steps in the overall tour path. The running times for the tour planner and the camera planner are 0.19 and 0.38 seconds, respectively. The tour guide's animation is generated on demand. The average running time for the humanoid simulator to generate one walking step is about 3 ms, which is negligible during the animation.

In summary, the overall planning time of the system is quite satisfactory for our interactive application. After the user selects locations of interests and clicks on the "go" button, complete paths for the tour guide and the virtual camera can be ready for execution in fractions of a second to a few seconds.

## 7. Conclusion and Future Extensions

Although networked virtual reality is a promising direction for realistic experiences in a virtual world, there are several problems to overcome in order to make 3D navigation a more pleasant experience. We believe that the key will lie on using machine intelligence to bring up the level of control that a user needs to provide. In this paper, we described an implemented system capable of providing customizable guided tours to a remote user. This system incorporates motion-planning algorithms to generate collision-free paths for the tour guide and for the virtual camera. The motions for the tour guide and the camera are then animated in a 3D VRML browser. Experiments have shown that the generated guided tours including the human animations are all very effective, and the computation times for each of the planning or simulation modules are adequate for on-line applications.

In the current system, the only goal of the tour planner is to pass through all locations specified by the user. However, the purpose of taking a guided tour in a virtual world might be more than being there. For example, the tour guide may stop at a popular tourist spot and provide the user with richer contents in forms of body gestures or narration. The user may also choose a different tracking strategy that allows the tour guide to leave its field of view at different occasions for certain amount of time. We are in the process of capturing these Cinematographical idioms in a finite state machine and allowing them to be incorporated into the tour path[8].

Although the traversing sequence found with our greedy strategy are all quite satisfactory, an optimal sequence is still desirable if it can be found in a short amount of time. We are in the process of designing an algorithm taking advantages of the topological characteristics of the skeleton in an architectural environment. We observe that the structure of the skeleton in a typical floor map is not a dense graph. Instead, it is more like a tree structure except for small local cyclic loops. If we can isolate and process these loops separately, an optimal solution can be found by composing the optimal solutions from these general but smaller TSP problems.

## Acknowledgements

This work was partially supported by grants from National Science Council under NSC 88-2218-E-004-002.

## References

- [1] J. Barraquand and J.-C. Latombe, "Robot Motion Planning: A Distributed Representation Approach," *The International Journal of Robotics Research*, 10(6), pp. 628-649, December 1991.
- [2] C. Becker, H. Gonzalez-Banos, J.-C. Latombe, and C. Tomasi, "An Intelligent Observer," in *Proceedings of International Symposium on Experimental Robotics*, pp. 94-99, 1995.
- [3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, "Introduction to Algorithms," The MIT press, pp.969-974, 1994.
- [4] P. Doyle and B. Heyes-Roth, "An Intelligent Guide for Virtual Environments," *Proceedings of the First International Conference on Autonomous Agents*, pp. 508-509, February 1997.
- [5] M. Erdmann and T. Lozano-Perez, "On Multiple Moving Objects," AI Memo No. 883, Artificial Intelligence Laboratory, MIT, 1986.
- [6] M. Girard and A.A. Maciejewski, "Computational Modeling for the Computer Animation of Legged Figures," *Computer Graphics, (Proceedings of SIGGRAPH95)*, 19(3), pp. 263-270, 1985.
- [7] H. H. Gonzalez-Banos, L. Guibas, J.-C. Latombe, S.M. LaValle, D. Lin, R. Motwani, and C. Tomasi, "Motion Planning with Visibility Constraints: Building Autonomous Observers," in *Proceedings of the Eighth International Symposium of Robotics Research*, Hayama, Japan, October 3-7, 1997.
- [8] L.-W. He, M. Cohen and D. Salesin, "The Virtual Cinematographer: a Paradigm for Automatic Real-time Camera Control and Directing," *Computer Graphics Proceedings, SIGGRAPH95*, pp. 217-224, 1995.
- [9] J.-C. Latombe, "Robot Motion Planning," Kluwer Publishing, pp. 318-334, 1990.
- [10] J.P. Laumond, "Obstacle Growing in a Non-Polygonal World," *Information Processing Letters*, 25(1), pp. 41-50, 1987.
- [11] S. M. LaValle, H. H. Gonzalez-Banos, C. Becker, J.-C. Latombe, "Motion Strategies for Maintaining Visibility of a Moving Target," in *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, 1997.
- [12] T.-Y. Li, L.-K. Gan, and C.-F. Su, "Generating Customizable Guided Tours for Networked Virtual Environments," in *Proceedings of 1997 National Computer Symposium (NCS'97)*, Taichung, pp. D189-D195, December 1997.
- [13] T.-Y. Li and J.-C. Latombe, "Online Manipulation Planning for Two Robot Arms in a Dynamic Environment," in *International Journal of Robotics Research*, 16(2):144-167, April, 1997.
- [14] T.-Y. Li and T.-H. Yu, "Planning Object Tracking Motions," in *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, May 1999.

- [15] T. Lozano-Perez, "Spatial Planning: A Configuration Space Approach," *IEEE Transaction on Computers*, 32(3), pp108-120, 1983.
- [16] Chris Marrin, "Proposal for a VRML 2.0 Informative Annex: External Authoring Interface Reference," URL:<http://www.vrml.org/WorkingGroups/vrml-eai/ExternalInterface.html>, Nov. 21, 1997.
- [17] J.H. Reif, "Complexity of the Mover's Problem and Generalizations," *Proceedings of the 20th IEEE Symposium on Foundations of Computer Science*, pp. 421-427, 1979.
- [18] J. Rickel and W.L. Johnson, "Integrating Pedagogical Capabilities in a Virtual Environment Agent," *Proceedings of the First International Conference on Autonomous Agents*, February 1997.
- [19] C. Rose, B. Guenter, B. Bodenheimer, and M. Cohen, "Efficient Generation of Motion Transitions using Spacetime Constraints," *Computer Graphics Proceedings, SIGGRAPH96*, pp. 147-154, 1996.
- [20] K. Sims and D. Zeltzer, "A Figure Editor and Gait Controller for Task Level Animation," *SIGGRAPH Course Notes*, 4, pp. 164-181, 1988.
- [21] "Specification for a Standard VRML Humanoid," <http://ece.uwaterloo.ca/~h-anim>, September 10, 1998.
- [22] M. Unuma, K. Anjyo, and R. Takeuchi, "Fourier Principles for Emotion-Based Human Figure Animation," *Computer Graphics Proceedings, SIGGRAPH95*, pp. 91-96, 1995.
- [23] "VRML97 Specification," International Standard ISO/IEC 14772-1:1997, <http://www.vrml.org/Specifications/VRML97>.
- [24] A. Witkin and Z. Popovic, "Motion Warping," *Computer Graphics Proceedings, SIGGRAPH95*, pp. 105-108, 1995.

