Joint EUROGRAPHICS - IEEE TCVG Symposium on Visualization (2002)
D. Ebert, P. Brunet, I. Navazo (Editors)

Viewpoint Entropy: A New Tool for Obtaining Good Views of Molecules

Pere-Pau Vázquez† Miquel Feixas‡ Mateu Sbert‡ Antoni Llobet⋆

† Departament de Llenguatges i Sistemes Informàtics, Universitat Politècnica de Catalunya, Barcelona, Spain
‡ Institut d'Informàtica i Aplicacions, Universitat de Girona, Girona, Spain
⋆ Departament de Química, Universitat de Girona, Girona, Spain

Abstract

The computation of good viewpoints is important in several fields: computer graphics, removal of degeneracies in computational geometry, robotics, graph drawing, etc. However, in areas such as computer graphics there is no consensus on what a good viewpoint means and, consequently, each author uses his or her own definition according to the requirements of the application. In this paper we present a formal measure strongly based on Information Theory, viewpoint entropy*, that can be applied to certain problems of Computer Graphics such as automatic exploration of objects or scenes and Scene Understanding. We also define a new measure, the* orthogonal frustum entropy*, in order to fulfill the requirements needed to visualize molecules. We design an algorithm that makes use of graphics hardware to accelerate computation, and whose complexity depends mainly on the number of views we want to analyze. Computation of good views of molecules is useful for molecular scientists, a field which includes practitioners from Crystallography, Chemistry, and Biology.*

1. Introduction

In the scientific world we often deal with three dimensional data. Unfortunately, in most cases, we only have at our disposal 2D media (e. g. images) to inspect it. Therefore it is frequently desirable to obtain 2D representations that allow us to understand the elements being studied. This problem appears in several fields: computer graphics, graph drawing, data visualization, knot theory, robotics, etc. Despite that, the *goodness* of a 2D representation (a view or image) can be differently evaluated depending on the problem we have in mind. As an example, certain views which are good for Image-Based Modeling purposes can be useless for object recognition.

In this paper we present a measure, the *viewpoint entropy*, that can be seen as the amount of information of a scene that can be captured from a point. Viewpoint entropy has been used in a previous work to address some problems in computer graphics, such as Scene Understanding or automatic exploration of objects or scenes[16]. Scene understanding techniques[2] deal with the problem of selecting a set of images that are enough to make the user understand the scene being represented. We will modify viewpoint entropy in order to cope with orthogonal projections (the kind of projections used by molecular science practitioners, who belong to several fields, such as Crystallography, Chemistry, and Biology) and see how this measure can be applied with molecules. A hardware-based algorithm which computes a solution with user-defined approximation is then presented.

The rest of the paper is organized as follows: In Section 2 we study the meaning of the term *good view* in different fields and present the main features of the problem we address: molecular visualization. Section 3 presents the measure of *viewpoint entropy* and introduces two new measures to deal with images that do not cover 360 degrees, the latter specially adapted for the case of molecular models. In Section 4 we present our hardware-based algorithm and show the results, and finally, in Section 5 we discuss our achievements and point out some lines of future work.

2. Previous Work

Scientific practitioners frequently need to visualize 3D data. Generally, only a 2D displaying method, e. g. pictures or a computer screen, is available. This poses some difficulties in

recognizing or understanding the objects being analyzed as the projected 2D images may have different topology to the 3D counterparts, or might not show enough information of the scene or object. On the other hand, certain 2D projections have some characteristics that make them useful for inspecting 3D objects, which are called either *nice*, or *regular*, or simply *good* projections. There are a number of different criteria that describe this notion of *good view* depending on the application in mind.

Automatic selection of good viewpoints has recently started to receive great attention in Computer Graphics due to the emergence of Image-Based Rendering[12, 6, 10] techniques. As the data is volumetric, then the occlusion problems that arise are harder to solve than in computational geometry[15].

Kamada and Kawai[9] and Roberts and Marshall[13] consider a direction to be good if it minimizes the number of degenerated faces when the scene is projected orthogonally. However, this condition is not sufficient because it fails when comparing scenes with equal number of degenerated faces and it does not ensure that the user will see a large amount of detail, as discussed in[2]. Barral et al[2] and Dorme[5] modify Kamada's coefficient in order to cope with perspective projections. Then they create a heuristic with some other parameters that weigh both the number of faces seen from each point and the projected area, moreover they add an exploration parameter which accounts for the faces already visited. However, they admit that they have not been able to determine a good weighting scheme for the different factors. This causes some problems with objects containing holes, as these are not captured properly by the algorithm. Vázquez *et al*[16] present an Information Theory-based measure called *viewpoint entropy* which can be used to determine which points of view in a scene show higher information.

In robotics literature, the goal of selecting a small set of cameras which allow us to observe the whole object has also been studied under the name of sensor planning or next best view selection. Different assumptions are made in next best view systems to simplify the problem. Several systems require a CAD model of the scene to be known a priori. The two main approaches are: search-based and silhouette-based. Search-based methods use optimization criteria to search a group of potential viewpoints of the next best view. Many of these methods employ range images to carve away voxels in a volumetric space. Wong *et al*[17] present an algorithm that searches all possible viewpoints, and selects the next best view as the one that can carve away the most empty space voxels. This system is effective, but as pointed out by Massios and Fisher[11], such an approach may result in views that observe surfaces at very oblique angles. Other approaches use the silhouettes of objects. For example, Abidi[1] develops a method that employs information theory. For a given view, a silhouette is divided into segments of equal lengths. Then, an information measure that computes the geometric and photometric entropy is found for each segment. The segment with the minimal entropy is chosen to select the next best view. This method assumes that by moving the camera to observe the segment containing the least information better, more information about the scene will be captured. Silhouette-based methods can often compute next best views more quickly than search-based approaches. However, it is not always possible to generate an accurate silhouette in an image for an arbitrary (for example indoor) scene.

There are few papers which refer to the viewpoint selection process for Image-Based Modeling, however, most papers select a fixed set of cameras placed in arbitrary positions[7, 12, 10]. Stürzlinger[14] creates a method for sampling all visible surfaces but does not address the problem of adequate coverage. Fleishman *et al*[6] present an algorithm that adequately samples the surfaces visible from a certain walking region by placing the camera on a large number of positions on the boundary of the walking zone. The coverage quality criterion for a polygon is based on its projected area on a hemisphere for a camera position. The set of cameras is selected by choosing the cameras that sample a higher number of polygons at an appropriate rate, which does not ensure that the amount of information they provide from the scene is high. If we had a scene with certain regions covered with a lot of very small polygons, this method could first sample parts of the scene that cover small areas instead of choosing other regions which cover larger portions of an image with less (or closer) polygons[2, 5, 16]. Hlavac et al[8] use a set of images to represent an object. They choose a set of reference images positioned around the object in intervals that guarantee error bounds below some threshold when interpolating intermediate views. This method is not intended to measure the quality of a view, as each image is compared with the previous one and only chosen if the degree of dissimilarity is high enough. So it is not possible to evaluate the goodness of completely different projections from the same scene.

Knot theory, graph drawing, and computational geometry usually deal with ideal objects consisting of segments and points[15]. This is not suitable for molecules as the volume of the different atoms encodes its actual covalent radius.

2.1. Molecular Visualization

Molecular scientists are concerned with obtaining good views of molecules, which allow them to infer their chemical or physical properties. In this case, we have three conditions:

1. Molecules are rigid objects.
2. Atoms (vertices) and bonds (edges) have a volume, they are therefore represented as spheres and cylinders respectively.
3. Information is encoded in the colour of the atoms and the size of their radius.

Unlike computational geometry approaches, where a good view is enough, we seek for an optimal projection, as

we want to see all the elements of a molecule under the better projection. The good viewpoints of a molecule consist both of the points where we see most of the atoms and bonds and the ones which show how the molecule is ordered in space. We will show in Section 3.3 that they correspond to the points with maximum and minimum orthogonal entropy (the ones which provide maximum and minimum amount of information respectively).

3. New Measures

In this section we present an Information Theory[3, 4] measure, viewpoint entropy. It was first introduced by Vázquez *et al*[16] and can be interpreted as the amount of information of a scene which can be seen from a point. Viewpoint entropy takes into account the number of faces and the projected area in a virtual scene. Our concept of information is based on *visibility*.

3.1. Viewpoint Entropy

Viewpoint entropy is based on an Information Theory measure, the Shannon Entropy[3, 4]. The Shannon Entropy of a discrete random variable X with values in the set $\{a_1, a_2, ..., a_n\}$ is defined as

$$H(X) = -\sum_{i=1}^{n} p_i \log p_i$$

where $p_i = Pr[X = a_i]$, the logarithms are taken in base 2 and when $p_i = 0$, $p_i \log p_i$ is 0 for continuity reasons[3, 4]. As $-\log p_i$ represents the *information* associated with the result a_i, the entropy gives the *average information* or the *uncertainty* of a random variable. The unit of information is called a *bit*. Let scene S consist of a set of N_f faces. We are going to use as probability distribution the relative area of the projected faces over the sphere of directions centered in the viewpoint p, as in Figure 1. Thus, the *viewpoint entropy* of a point p from a scene S is defined as[16]:

$$I(S, p) = -\sum_{i=0}^{N_f} \frac{A_i}{4\Pi} \log \frac{A_i}{4\Pi} \qquad (1)$$

where A_i is the projected area of face i and 4Π is the solid angle of the sphere. Hence, $A_i/4\Pi$ represents the *visibility* of face i with respect to the point p. When $i = 0$, the area projected is the background. This is needed to have a well-built probability distribution function (cf. Vázquez *et al*[16]).

3.2. Perspective frustum entropy and orthogonal frustum entropy

In many cases what we really want to measure is the amount of information provided from a single image that does not cover all the sphere of directions, as this is the way we usually obtain the 2D representations. In order to do this we must consider two cases, the one where we have a perspective projection (either coming from real world or from a virtual scene), and the one where we deal with orthographically projected images (which are equivalent to a perspective projection with the point of view placed at infinity). To obtain such measures, we use the viewpoint entropy as a basis. For the perspective case, the measure is the same and we only have to change the total solid angle of the sphere to the total area of the intersection of the frustum pyramid with the sphere. For the orthogonal case we have the same formula but the pixels must not be weighted by the angle they subtend. Thus, to measure the *orthogonal frustum entropy* of a view we can apply the following formula:

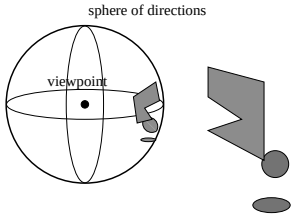

Figure 1: *Computation of the viewpoint entropy by projecting the objects onto a bounding sphere of the viewing point.*

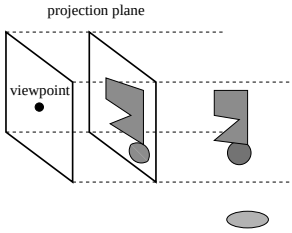

Figure 2: *Only the objects inside the frustum volume are considered for the entropy computation.*

$$I_O(S, p) = -\sum_{i=0}^{N_f} \frac{N_{pix_i}}{N_{pix_F}} * \log \frac{N_{pix_i}}{N_{pix_F}}, \qquad (2)$$

where N_{pix_i} is the number of the projected pixels of face i, and N_{pix_F} is the total number of pixels of the image. This measure is appearance-based in the sense that it only measures what we can really see. This means that we will apply formula (2) to the objects that project at least one pixel on the screen, thus, perceivable by an observer.

3.3. Good viewpoints of a molecule

In this section we analyze the requirements of molecular visualization and see how they can be fulfilled using the orthogonal viewpoint entropy measure.

A good viewpoint can be defined as the one that gives more information about the scene (Vázquez *et al*[16]). Two cases are specially important for molecular scientists:

1. Projections with high orthogonal entropy of single molecules: they give a lot of information about the molecules being observed. Thus, these projections allow us to see most of the atoms composing the molecule and

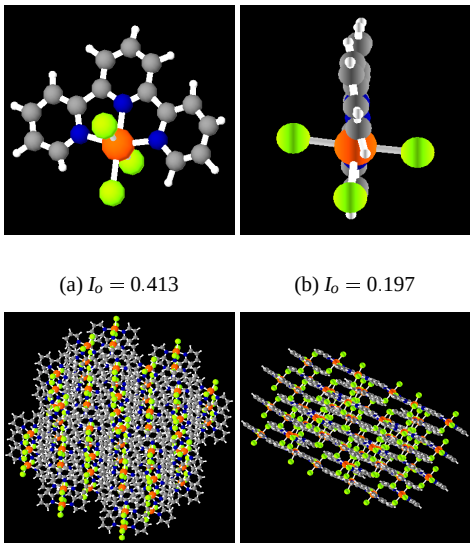

(a) $I_o = 0.413$ (b) $I_o = 0.197$

(c) $I_o = 4.5550$ (d) $I_o = 2.0805$

Figure 3: *Figures (a) and (c) show the points of maximum orthogonal entropy of a molecular representation of Ru(III) trichloro-2,2':6',2"-terpyridine and an arrangement of the same molecule. Figures (b) and (d) show the respective points of minimum orthogonal viewpoint entropy.*

the distances and angles of the bonds. These features are relevant for molecular scientists because they help to infer chemical properties of the molecule.

2. Projections with low orthogonal entropy of arrangements of the same molecule: The ordering in space determines physical properties of the molecules, which can then be discovered by scientists from these views.

A single projection of maximum entropy might not be enough to represent all atoms and bonds. In Section 4 we present a hardware-based algorithm that finds the minimum set of images that shows all these features. The best image is the one which maximizes the information provided. In Figure 3*a* we can see a representation of a Rutenium compound seen from the position of maximum entropy, and in 3*b* seen from the point of minimum viewpoint entropy. Note that in 3*a* all the components of the molecule are visible. The respective arrangements of several of these molecules appear in 3*c* and 3*d*. Note that in 3*c* it is easy to perceive a regular ordering in some direction.

4. Hardware-based orthogonal entropy computation

With today's technology, powerful graphics cards are affordable for low-end personal computers. They permit the design of algorithms using OpenGL and, consequently, rendering can be done in real time. In this section we show a hardware-accelerated algorithm which enables us to compute the orthogonal entropy of an image in milliseconds without major optimizations.

```
Select a set of points placed in regular positions all around the object
for all the points do
    Compute the orthogonal entropy and store it
    Store a bitmap encoding the visibility of the faces from the point
end for
Order the points in decreasing viewpoint entropy
Select the first point      {the one with maximum orthogonal entropy}
Accumulate the visited faces in a bitmap
i ← 0
while i < totalPoints and not finished do
    if numFacesNotSeen(i) > threshold then
        Select point i
        Accumulate the visited faces in a bitmap
        finished ← isFinished(numberofVisitedFaces)
    end if
    i ← i + 1
end while
```

Figure 4: *Algorithm that computes the set of view with high entropy of a molecule and which covers all the elements.*

4.1. Data representation

We have a model composed by spheres and cylinders, and in equations (1) and (2) the elements we project are faces. We usually assume that we have scenes composed by polygons, and elements such as spheres or cylinders are discretized. This is due to the fact that the usual scenes can have different optical properties for different polygons, and thus we are interested in capturing the information of *the whole* scene. In Molecular Visualization, atoms and bonds look the same from all points of view. Consequently, when we see an atom we know how it will look from the back side. That's why we do not need to discretize the scene. Equivalent results are obtained if we discretize in equal sized triangles for same sized objects (otherwise it would lead to errors as the orthogonal entropy of two faces discretized in different manners has also different values). The finer the discretization the better the results would be.

We assume the molecule is centered at point $(0,0,0)$. To compute the best viewpoints we use the algorithm that appears in Figure 4. Orthogonal entropy is computed by colour-coding the faces in an item buffer. The images are captured by a camera placed in the current position and pointing to the origin. Then, orthogonal entropy is calculated by summing up all the pixels with the same colour (i.e. all the pixels coming from the same face) and applying (2). Then, the viewpoints are ordered in decreasing entropy and a loop selects the points with high entropy which see a certain number of faces (above a threshold) not yet visited. The threshold can depend on the number of not yet visited faces, to ensure we cover all the faces and to avoid adding views which show only one new atom or bond.

4.2. Results

We have evaluated our algorithm with different molecules. Molecular scientists need to visualize first the molecule all alone and then a group of them ordered in space. Two kinds of viewings are interesting, the ones that show a lot of details on the molecule and the ones which allow us to see how it is ordered in space, these correspond to the ones which present low entropy. Consequently, our implementation searches the points of high orthogonal entropy and also the points of low orthogonal entropy. In Figure 5*a* we see a molecular representation of a tetrabromic salt from the point of maximum orthogonal entropy. This view allows to see that this crystal packing generates channels between molecules. Their proximity permits to infer the presence of intermolecular hydrogen bridges. Although higher resolution would obtain slightly better images (with the back atoms completely occluded) this is not a severe drawback as the ones obtained are already good enough for the users, as they also enable us to see that there are some more atoms in the back side. Figure 5*b* shows the same molecule from the minimum entropy view. Figures 5*c* and 5*d* show the points of maximum and minimum entropy for an arrangement of the same molecule. In Figure 6 two forms of Carbon are shown. Figure 6*a* shows a graphite and Figure 6*b* a diamond. Note that the arrangement into layers of graphite may indicate that there is a direction of easy exfoliation of the molecule (as it really occurs) while the structure of diamond with bonds in different directions of space make it a strong molecule.

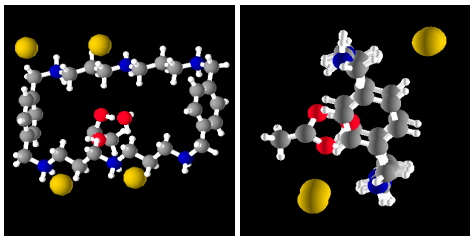

(a) $I_o = 0.7844$ (b) $I_o = 0.4582$

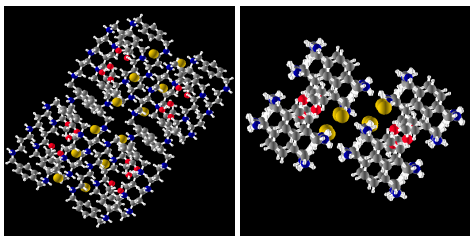

(c) $I_o = 2.593$ (d) $I_o = 1.104$

Figure 5: *Figures (a) and (b) show a molecular representation of a tetrabromic salt of a hexaazamacrocyclic ligand seen from the points of maximum and minimum entropy respectively. Figures (c) and (d) show the views of maximum and minimum entropy for a set of these molecules.*

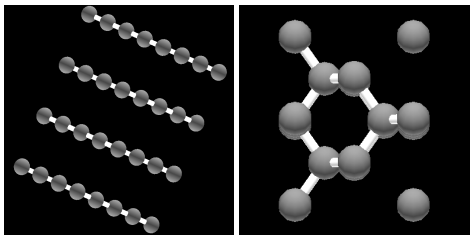

(a) (b)

Figure 6: *Figures (a) and (b) show the minimum entropy views of the molecular representations of two Carbon forms, graphite and diamond respectively. From these views molecular scientists can infer the resistance to physical pressure. While diamond is very strong due to the bonds in three directions, graphite's layered structure makes the molecule easily exfoliable.*

4.3. Performance

The complexity of our approach depends mainly on the number of views being evaluated and the image size. It depends only weakly on the number of faces of the model, and models with several hundreds of faces behave like models with several thousands. Our algorithm has several advantages:

- The number of views selected for evaluation can be chosen by the user.
- The number of resulting views can be also set by the user.
- Adaptive methods can be applied to find better results.

As the entropy is defined as a continuous function, adaptive methods could be used to improve the results produced by a sparsely sampled first evaluation. For a resolution of 400×400 pixels, we obtain a performance of 17-18 fps on a AMD-K7 processor at 700 MHz. We used a Riva TNT2 graphics card and the algorithm worked in the front buffer, and no optimizations are made. Several improvements can be applied if we are seeking for better frame rates, such as the use of bounding boxes to accelerate image analysis, visibility preprocess, etc.

The process of synthesizing a molecule is very complex. It consists of two main steps: Synthesis and purification of the molecule, and the characterization step where an X-ray diffraction process generates the crystallographic data. This is the model we visualize. The whole process can take several years, but the characterization step can be obtained in a period which varies from several hours up to more than one week. When the model has been created, our algorithm can compute the optimal views in a couple of minutes, so if we add our method to the characterization pipeline its cost is negligible. As an automated process did not exist, scientists had to spend up to several hours to obtain a good view with the common tools.

5. Conclusions and Future Work

Visualization of molecules is relevant for molecular science, a discipline which falls in several areas such as Crystallography, Chemistry, and Biology. Obtaining such views is time consuming for molecular scientists. In this paper we have presented a new approach that computes automatically good viewpoints of a molecule. First, we have introduced a new measure, the *orthogonal viewpoint entropy*. Then we have designed a hardware-based algorithm that automatically obtains good views in a couple of minutes using this measure. The results obtained by the algorithm are a set of views which fulfill the requirements of molecular visualization: they reveal most of the detail of the molecule, and show how it is ordered in the space. The user can set the level of approximation of the algorithm, as well as the number of resulting images he or she wants to obtain. In most cases the views generated by our application can completely replace human involvement, otherwise, for highly complex compounds, they are a good starting point. It also provides several views that could be missed by scientists and it has proven to greatly simplify the chemist work.

In the future we will study the particular requirements of the visualization of DNA models, as they have several similarities with the problem addressed in this paper. Besides, we will study adaptive methods to compute orthogonal entropy. Moreover, some improvements that have been suggested are the computation of rotations around certain axes and the visualization of bonding angles.

Acknowledgments

This work was partially supported by BR98/1003 grant of Universitat de Girona, TIC2001-2416 from the Spanish government, and SGR2001-00296 from Catalan government.

References

1. B. Abidi. Automatic sensor placement. In *Proc. Intelligent Robots and Computer Vision: Algorithms, Techniques, Active Vision, and Materials Handling*, pages 387–398, 1995.
2. P. Barral, G. Dorme, and D. Plemenos. Scene understanding techniques using a virtual camera. In A. de Sousa and J.C. Torres, editors, *Proc. Eurographics'00, short presentations*, 2000.
3. R.E. Blahut. *Principles and Practice of Information Theory*. Addison-Wesley, 1987.
4. T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 1991.
5. G. Dorme. *Study and implementation of 3D scenes comprehension techniques.* PhD thesis, 2001.
6. S. Fleishman, D. Cohen-Or, and D. Lischinski. Automatic camera placement for image-based modeling. *Computer Graphics Forum*, 19(2):101–110, Jun 2000.
7. J.P. Grossman and William J. Dally. Point sample rendering. In George Drettakis and Nelson Max editors, editors, *Rendering Techniques'98*, pages 181–192. Springer-Verlag, 1998.
8. V. Hlavac, A. Leonardis, and T. Werner. Automatic selection of reference views for image-based scene representations. In *Lecture Notes in Computer Science*, pages 526–535, New York, NY, 1996. Springer Verlag. Proc. of European Conference on Computer Vision '96 (ECCV '96).
9. T. Kamada and S. Kawai. A simple method for computing general position in displaying three-dimensional objects. *Computer Vision, Graphics, and Image Processing*, 41(1):43–56, January 1988.
10. D. Lischinski and A. Rappoport. Image-based rendering for non-diffuse synthetic scenes. In George Drettakis and Nelson Max editors, editors, *Rendering Techniques'98*, pages 301–314, 1998.
11. N.A. Massios and R.B. Fisher. A best next view selection algorithm incorporating a quality criterion. In *Proc.of the British Machine Vision Conference*, 1998.
12. L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. *Proc. of SIGGRAPH 95*, pages 39–46, August 1995.
13. D.R. Roberts and A.D. Marshall. Viewpoint selection for complete surface coverage of three dimensional objects. In *Proc.of the British Machine Vision Conference*, 1998.
14. W. Stuerzlinger. Imaging all visible surfaces. In I. Scott MacKenzie and James Stewart, editors, *Proc. of the Conference on Graphics Interface (GI-99*, pages 115–122, Toronto, Ontario, June 2–4 1999. CIPS.
15. Godfried T. Toussaint. The complexity of computing nice viewpoints of objects in space. In *Vision Geometry IX, Proc. SPIE International Symposium on Optical Science and Technology*, San Diego, California, 2000.
16. Pere-Pau Vázquez, Miquel Feixas, Mateu Sbert, and Wolfgang Heidrich. Viewpoint selection using viewpoint entropy. In T.Ertl, B. Girod, G.Greiner, H. Niemann, and H.-P. Seidel, editors, *Vision, Modeling, and Visualization 2001*, 2001.
17. L. Wong, C. Dumont, and M. Abidi. Next best view system in a 3-d object modeling task. In *Proc. International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, pages 306–311, 1999.