Chapter 5

Behavioral Difference and Diversity

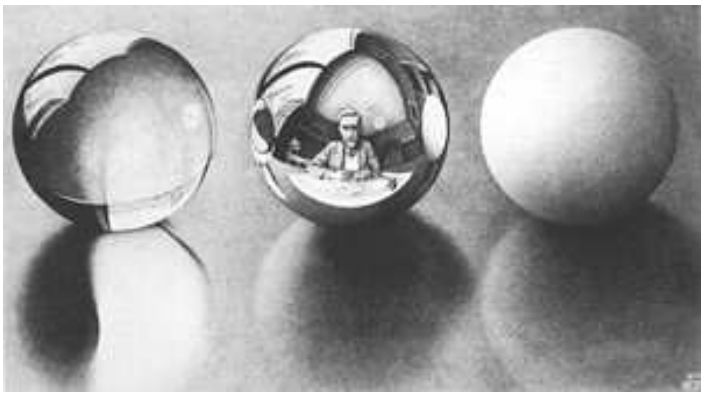

As research expands in multiagent intelligent systems, investigators need new tools for evaluating the artificial societies they study. It is impossible for example, to correlate heterogeneity with performance in multiagent robotics without a quantitative metric of diversity. Currently diversity is evaluated on a bipolar scale with systems classified as either *heterogeneous* or *homogeneous*, depending on whether agents differ [FM97, GM97, Par94]. Unfortunately, this labeling doesn't tell us much about the *extent* of diversity in heterogeneous teams. How can it be determined if one system is more or less diverse than another? Heterogeneity is better viewed on a sliding scale providing for quantitative comparisons. Such a metric would enable the investigation of issues like the impact of diversity on performance, and conversely, the impact of other task factors on diversity.

Diversity may not always be desirable. In fact, experimental results presented

later in this dissertation show that for some multi-robot tasks homogeneous robot teams perform better than diverse teams. The aim of this work is to discover when diversity is important and which conditions give rise to it in learning teams. An objective quantitative metric is required for a principled investigation of these issues.

We focus specifically on diversity in teams of mechanically similar agents that use reinforcement learning to develop behavioral policies. Evaluation of diversity in teams of mechanically similar robots is challenging because when agents differ, they differ only in their behavior. Behavior is an especially interesting dimension of diversity in learning systems since as they learn, agents effectively choose between a hetero- or homogeneous society. The metrics introduced here will help researchers investigate the origin and benefits of diversity in these learning systems.

Measurement of diversity is really a three-pronged problem: measurement of individual agent *difference*, agent *classification* (clustering/clumping) based on inter-agent differences, and overall societal *diversity* based on the classification. Each of these topics is covered in a separate section below.

Social entropy, inspired by Shannon's information entropy [Sha49], is presented as an appropriate measure of diversity in robot teams. It captures important components of the meaning of diversity, including the number and size of groups in a society. In order to evaluate the diversity of a team, however, a way to categorize or differentiate the behavior of individuals is also required. To address this, a measure of *behavioral difference* that provides for agent categorization is presented in later sections. Difference refers to disparity between two specific agents, while diversity is a measure of the entire society. Finally, the utility of the metrics is demonstrated in several example applications, including a detailed evaluation of a foraging robot team.

5.1 Diversity

We begin with an examination of the meaning of diversity and the challenges it presents to measurement.

What does *diverse* mean? Webster [MW89] provides the following definition:

di.verse *adj* **1:** differing from one another: unlike **2:** composed of distinct or unlike elements or qualities

Clearly, difference plays a key role in the meaning diversity. In fact, an important challenge in evaluating robot societal diversity is determining whether agents are alike or unlike. Assume for now that any two agents are either alike (in the same behavioral group) or not. It may be true that the degree of difference *is* important but that issue is addressed later.

Consider what *diverse* means for societies composed of distinct groups. To make the discussion more concrete, suppose the "society" under examination is a collection of toy blocks of four different shapes: circles, squares, triangles and stars. Figures 5.1 and 5.2 illustrate several sets of blocks as examples of the different ways groupings can differ. The goal is to develop a quantitative metric that captures the meaning of diversity illustrated in these examples.

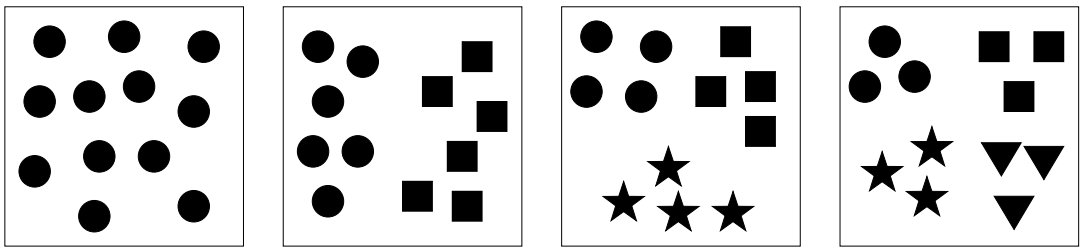

Figure 5.1: Four collections of toy blocks. The number of subsets in each group grows from one in **a** to four in **d**. Should measured diversity depend on the number of subsets?

First, how should the number of distinct groups in a society impact measured diversity? In Figure 5.1 we have four sets of 12 blocks. Each set has a different number of subsets of homogeneous shapes; from one subset in Figure 5.1a (all circles) to four subsets in Figure 5.1d. In each case the subsets are uniform: there are the same number of blocks in each subset. This example suggests that the number of subsets in a society is an important component of measured diversity.

Now consider Figure 5.2. Which group of blocks is more diverse? In both cases there are exactly 12 blocks and exactly two different types of block. In Figure 5.2a however, there is a much higher proportion of circles than in 5.2b where there is an equal number of circles and squares. This example suggests that the relative proportion of elements in each subset is important component of diversity.

The examples above highlight the fact that the *distribution* of agents into subsets of homogeneous elements is at the core of the meaning of diversity for multiagent societies. We make the following commitment: **the measured diversity of a**

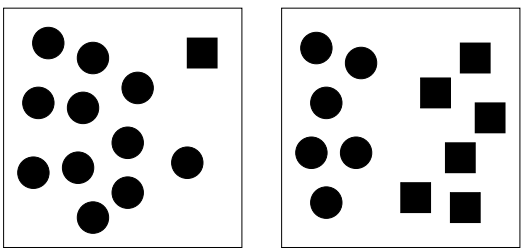

Figure 5.2: In both of these groups there are the same number of blocks and the same number of subsets, but the proportion of elements in each subset is different.

multiagent society depends on the number of subsets it contains and the proportion of agents in each subset.

5.1.1 Introducing social entropy

How should diversity be quantified? Shannon faced a similar problem when he sought to quantify the uncertainty, or randomness, of an information source [Sha49]. The uncertainty of an information source has important implications for communications systems, particularly with regard to the minimum bandwidth required to transmit error-free messages. Interestingly, the properties Shannon sought in a measure of information uncertainty are also important for a metric of societal diversity.

Shannon's solution, *information entropy*, is easily adapted to suit our needs in a diversity metric. The remainder of this section presents the mathematical basis for *social entropy* and explains why it is an appropriate measure of multiagent system diversity.

It is useful at this point to introduce the following notation:

- $\mathcal{R}$ is a society of N agents with $\mathcal{R} = \{r_1, r_2, r_3 ... r_N\}$
- $D(r_i, r_j)$ is the difference between agents r_i and r_j. Difference is used to classify agents into homogeneous subsets.
- $\mathcal{C}$ is a grouping of $\mathcal{R}$ into M possibly overlapping subsets.
- c_i is an individual subset of $\mathcal{C}$ with $\mathcal{C} = \{c_1, c_2, c_3 ... c_M\}$
- $p_i = \frac{|c_i|}{|\mathcal{C}|}$ is the proportion of agents in the ith subset; $\sum p_i = 1$.

Measured diversity is therefore a function of M and the p_is. Assume that a diversity function exists and call it H. The diversity of a society partitioned into M subsets is written $H(p_1, p_2, p_3, ..., p_M)$. So, for instance, the diversity of the group of

blocks depicted in Figure 5.2a is $H(\frac{1}{12}, \frac{11}{12})$, while the diversity for the group of blocks in Figure 5.2b is $H(\frac{1}{2}, \frac{1}{2})$. The diversity of a particular robot society $\mathcal{R}_a$ can also be expressed $H(\mathcal{R}_a)$.

Shannon prescribed three properties for a measure of information uncertainty. With slight changes in notation, they are equally appropriate for a measure of societal diversity:

Property 1 continuous: H should be continuous in the p_i.

Property 2 monotonic: If all the p_i are equal, $p_i = \frac{1}{M}$, then H should be a monotonic increasing function of M. In other words, if there are an equal number of agents in each group, more groups implies greater diversity.

Property 3 recursive: If a multiagent society is defined as the combination of several disjoint sub-societies, H for the new society should be the weighted sum of the individual values of H for the sub-groups. This property is important for the analysis of recursively composed societies (e.g. [MAC97]).

The meaning of the requirement that H be recursive is illustrated in Figure 5.3. The two groups on the left are combined into a new society on the right. The groups on the left have diversities $H(\frac{5}{6}, \frac{1}{6})$ (top) and $H(\frac{1}{2}, \frac{1}{2})$ (bottom). The diversity of the new 18 element society is $H(\frac{5}{18}, \frac{1}{18}, \frac{6}{18}, \frac{6}{18})$. Because the sub-groups contribute $\frac{1}{3}$ and $\frac{2}{3}$ of the agents to new society, the recursive criteria requires:

$$H(\frac{5}{18}, \frac{1}{18}, \frac{6}{18}, \frac{6}{18}) \quad = \quad H(\frac{1}{3}, \frac{2}{3}) + \frac{1}{3}H(\frac{5}{6}, \frac{1}{6}) + \frac{2}{3}H(\frac{1}{2}, \frac{1}{2})$$

In general, for a society $\mathcal{R}_c$ composed of two societies, $\mathcal{R}_a$ and $\mathcal{R}_b$, the recursive criteria ensures that:

$$H(\mathcal{R}_c) \quad = \quad H(\alpha, \beta)) + \alpha H(\mathcal{R}_a) + \beta H(\mathcal{R}_b)$$

where α is the proportion of agents in $\mathcal{R}_a$, β is the proportion of agents in $\mathcal{R}_b$ and $\alpha + \beta = 1$.

Shannon's *information entropy* meets all these criteria. The information entropy, $H(X)$ of a symbol system X is used in coding theory as a lower-bound on the average number of bits required per symbol to send multi-symbol messages. The random variable X assumes discrete values in the set $\{x_1, x_2, x_3 ... x_M\}$ (the alphabet to be encoded). The information entropy of X is given in bits as:

$$H(X) \quad = \quad -\sum_{i=1}^{M} p_i \log_2(p_i) \tag{5.1}$$

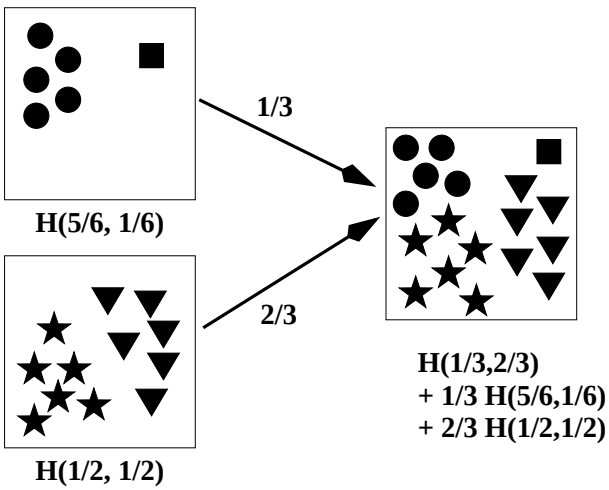

Figure 5.3: A new society (right) is generated by combining two others (left). The diversity of the new society is a weighted sum of the individual values of H for the subgroups.

where p_i represents the probability that $\{X = x_i\}$.

Equation 5.1 is adopted for the measurement of multiagent societal diversity. $H(\mathcal{R}_a)$ is the *simple social entropy* of agent society $\mathcal{R}_a$. Several example applications of this metric are given below.

Example evaluations

Consider the social entropy of a heterogeneous group composed of one square and three star shaped blocks. The society consists of four elements, $\mathcal{R} = \{r_1, r_2, r_3, r_4\}$. One element, r_4 (the square) is not equivalent to the others so there are two subsets, $C = \{c_1, c_2\}$, with $c_1 = \{r_1, r_2, r_3\}$ (the star class) and $c_2 = \{r_4\}$ (the square class). Then,

$$\begin{aligned} p_1 &= .75 \\ p_2 &= .25 \\ H(\mathcal{R}) &= -\sum_{i=1}^{2} p_i \log_2(p_i) \\ &= -((p_1 \log_2(p_1)) + (p_2 \log_2(p_2))) \\ &= -((.75 \log_2(.75)) + (.25 \log_2(.25))) \\ &= .811 \end{aligned}$$

The social entropy of this system is .811.

Next the social entropy of a homogeneous group is evaluated. The group consists of elements $\mathcal{R} = \{r_1, r_2, r_3, r_4\}$. Homogeneity implies there is only one class, so $\mathcal{C} = \{c_1\}$, and $c_1 = \{r_1, r_2, r_3, r_4\}$. Then:

$$\begin{aligned} p_1 &= 1 \\ H(\mathcal{R}) &= -\sum_{i=1}^{1} p_i \log_2(p_i) \\ &= -(p_1 \log_2(p_1)) \\ &= -(1 \log_2(1)) \\ &= 0 \end{aligned}$$

The entropy of a number of other example systems is given in Figure 5.4.

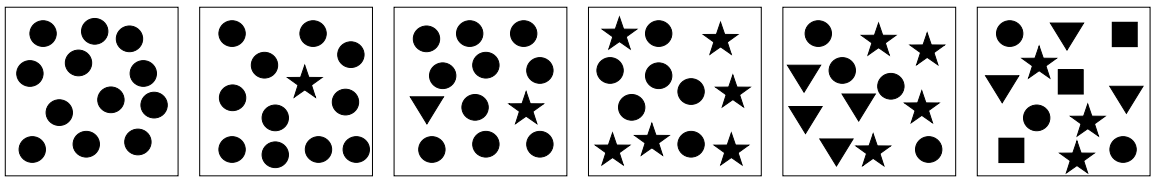

Figure 5.4: A spectrum of diversity. In the diagram above, each of the six squares encloses a multiagent system, from least diverse (homogeneous) on the left, to most diverse (maximally heterogeneous) on the right. The *social entropy*, a qualitative measure of diversity, is listed underneath each system.

5.1.2 Why entropy is a useful measure of diversity

In addition to Properties 1, 2 and 3, H has a number of additional properties that further substantiate it as an appropriate measure of diversity. First, as we would expect, H is minimized for homogeneous societies; these groups are the least diverse. Also, for heterogeneous groups H is maximized when there are an equal number of

agents in each subset. More precisely:

Property 4: $H = 0$ if and only if all the p_i but one are zero. In other words H is minimized when the system is homogeneous. Otherwise H is positive.

Property 5: For a given M (number of subsets), H is maximized when all the p_i are equal, i.e. $p_i = \frac{1}{M}$. This is the case when there are an equal number of agents in each group.

Property 6: Any change toward equalization of the proportions $p_1, p_2, \ldots, p_M$ increases H. Thus if $p_1 < p_2$ and we increase p_1, decreasing p_2 an equal amount so that they are more nearly equal, H increases. An important implication is that there are no locally isolated maxima.

Even if these properties are desirable in a diversity metric, why choose information entropy over any other function possessing the same properties? Because, as it turns out, information entropy (Equation 5.1) is the *only* function satisfying Properties 1, 2 and 3. Shannon proved this result using the mathematically equivalent properties he required of an information uncertainty metric [Sha49].

The utility of these properties for the measurement of diversity has drawn researchers in many other disciplines to adopt similar concepts of diversity. Several examples from the literature are included below.

Sociobiology: Wilson, the entomologist and creator of the field of sociobiology, includes this discussion in his book *The Diversity of Life* [Wil92]:

> Suppose that we encounter a fauna of butterflies consisting of 1 million individuals divided into 100 species. Say one of the species is extremely abundant, represented by 990,000 individuals, and each of the other species therefore comprises an average of about 100 individuals. One hundred species are present but, as we walk along the forest paths and across the fields, we encounter the abundant butterfly most of the time and each of the other species only rarely ... In a nearby locality we encounter a second butterfly fauna, comprising the same 100 species, but this time all are equally abundant, represented by 10,000 individuals each. This is a fauna of high equitability, in fact the highest possible. Intuitively we feel that **the high-equitability fauna is the more diverse of the two, since each butterfly encountered in turn is less predictable and therefore gives us more information on average.**

This view embraces the idea that societies with members equally distributed among sub-groups are the most diverse. It also suggests that diversity and information are closely related concepts.

Ecology: entropy is used by ecologists as a means of evaluating species diversity and for comparing diversity in differing environments [LVW83, LW80, Mag88]. Consider this passage from Magurran's book *Ecological Diversity and Its Measurement* [Mag88]:

> **The most widely used measures of diversity are the information theory indices.** These indices are based on the rationale that the diversity, or information, in a natural system can be measured in a similar way to the information contained in a code or message. It is calculated from the equation:
>
> $$H = -\sum p_i \log(p_i)$$
>
> where the quantity p_i is the proportion of individuals found in the ith species.

Social Science: information theoretic models of societal evolution are developed in [Bai90, May98] and others. Of these, Bailey's *Social Entropy Theory* [Bai90] is the most widely known. In his model, individuals are classified according to characters such as race, income and religion and the diversity of the society is calculated using information entropy.

Biological Classification: information entropy is an important tool for taxonomists as a mechanism for evaluating classification methodologies [SS73, JS71]. Classification trees are organized to maximize their information content.

Mathematics and Statistics: information entropy is suggested as a measure of diversity by a number of mathematicians and statisticians [Bev90, PPM92, MGL92].

Genetics and Evolution: Demetrius argues in *The thermodynamics of evolution* that the evolutionary process of genetic mutation leads to greater population diversity [Dem92]. The diversity is modeled as an increase in system entropy.

5.1.3 Limitations and alternative measures

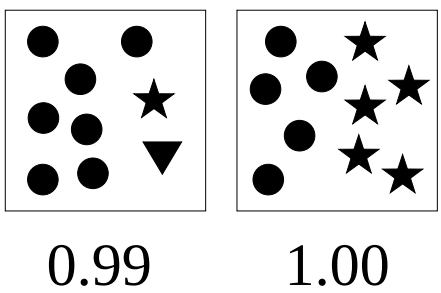

Figure 5.5: Two very different systems have similar entropy.

A potential limitation of social entropy as a diversity metric is the loss of information incurred when diversity is summarized in a single number. There are perhaps an

infinite number of societies matching any particular value of diversity. Figure 5.5 for example, illustrates two very different societies whose entropy differs by less than 0.01 bits.

A single number does not tell us how many classes of agents there are or how many agents in each class. This loss of information occurs whenever any characteristic of a multi-dimensional system is described in a single value. However such measurements are useful because they enable generalization and comparison. A thermometer, for example, does not reveal the position and velocity of every molecule in the environment but it does enable us to select our wardrobe for the day.

Another limitation of social entropy is that it lacks sensitivity to the degree of difference between agents (later in this chapter an augmented form of social entropy is introduced to address this weakness). Suppose, for example, we are evaluating the diversity of a number of agents distributed in a two-dimensional space (the dimensions may represent aspects of behavior or perhaps morphological axes). Agents that are close to one another are grouped in the same class. Figure 5.6 illustrates.

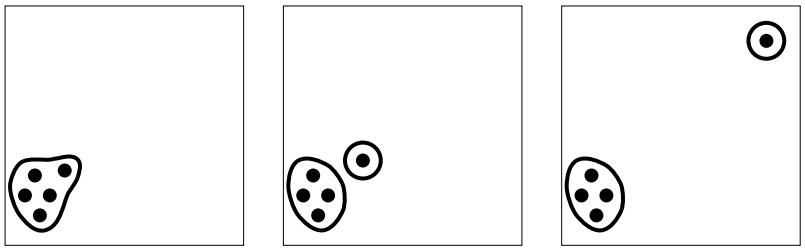

Figure 5.6: One difficulty in the analysis of diversity. Dots representing agents are plotted in a two-dimensional space. Lines enclose agents grouped in the same class. The entropy metric cannot distinguish between the systems illustrated in **b** and **c**.

The figure shows three systems. In each system, the four elements in the lower left remain unchanged, but from 5.6a to 5.6c a fifth agent appears in several locations progressively more distant from the others. In Figure 5.6a it is just close enough to be grouped with the others, while in 5.6b it is just far enough away to be placed in a separate category. The social entropy metric cannot differentiate between the distribution of agents in 5.6b and 5.6c because there is no difference in the number and size of the subsets. Also, the entropy measure finds a greater difference between the systems in 5.6a and 5.6b than between those in 5.6b and 5.6c.

One solution is to consider the maximum difference between agents as an addi-

tional component of diversity; e.g. the distance d in Figure 5.7. In the biological taxonomy literature d is referred to as *maximum taxonomic distance*. Taxonomic distance is useful, but as Figure 5.7 illustrates, it cannot serve as the only measure of diversity. This example shows two societies: one society with most of the agents grouped together, but one "outlier" at distance d (5.7a); and another society with two equally sized groups separated by the same distance (5.7b). Both of these systems have the same maximum difference but quite different distributions of agents into subgroups. Taxonomic difference captures the greatest difference between agents in the society but ignores the distribution of agents in the space. Again, this issue is addressed by an augmented version of social entropy presented later in the chapter.

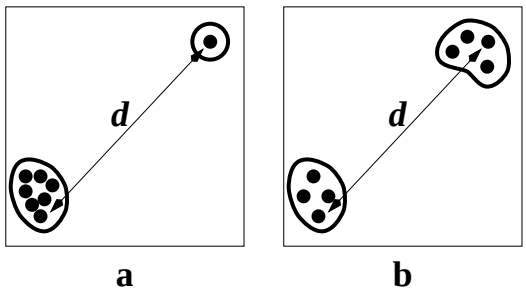

Figure 5.7: Maximum taxonomic distance is a useful metric, but it does not account for the distribution of elements in the space.

Meyer and McIntosh have developed an index of ethnic diversity used by *US News* and *USA Today* in stories concerning diversity issues [MM92]. Their index measures the probability that two people chosen at random (with replacement) will differ along at least one ethnic dimension. The index value ranges from 0 to 1, with greater diversity indicated by a larger value. A value of zero applies to a population in which everyone is the same. If every person is different from every other person on at least one dimension, the value is maximized. In practice the index can never reach unity because an infinite number of p_i would be required. The metric has intuitive appeal, and may be of interest as a measure of multiagent social diversity. Using the notation introduced above, Meyer's metric can be written:

$$H_m(X) \quad = \quad 1 - \sum_{i=1}^{M} p_i^2 \tag{5.2}$$

H_m shares mathematical properties 1, 2, 4, 5 and 6 with social diversity, but it does not provide for recursively defined societies (Property 3).

5.2 Classification and hierarchic entropy

The discussion of diversity left open the question of how agents are classified into groups. It was assumed that any two agents are either alike (in the same group) or unlike. In actuality, the robotic agents to be classified are distributed in a multi-dimensional space where the dimensions correspond to components of behavior and difference corresponds to the distance between agents in the space. Difference between agents is likely to vary continuously instead of in the binary manner assumed previously.

The limitations of social entropy discussed in Section 5.1.3 suggest that the diversity calculation would be improved if consideration were given to the spatial structure the system. Here "spatial structure" refers to the distribution of elements in the classification space. In other words, the clumpiness of the system and the distribution of the clumps in the space are important.

The challenge of finding and characterizing clumps or clusters of elements distributed in a continuous multi-dimensional space is exactly the problem faced by biologists in building and using taxonomic systems. In the case of biology the dimensions of the space represent aspects of morphology or behavior that distinguish one organism from another. In this research the dimensions are the components of agent behavior that distinguish one robot from another.

The aims of taxonomic classification are distinct from other types of classification in that one goal is to arrange the elements in a hierarchy reflecting their distribution in the classification space. Conversely, many classification tasks only require a simple partitioning of the space (e.g. categorizing e-mail into folders). Taxonomic trees are potentially more useful in the analysis of diversity than simple partitionings because they provide more information about the society's spatial structure.

Biology offers a rich literature addressing this problem. In fact, an entire field — *numerical taxonomy* — is devoted to ordering organisms hierarchically using principled numerical techniques [SS73, JS71]. Many of the approaches in numerical taxonomy are directly applicable to the problem of robot classification. They include mechanisms for building and analyzing classification structures (e.g. taxonomic

trees) and for identifying organisms on the basis of these structures.

5.2.1 Tools from numerical taxonomy

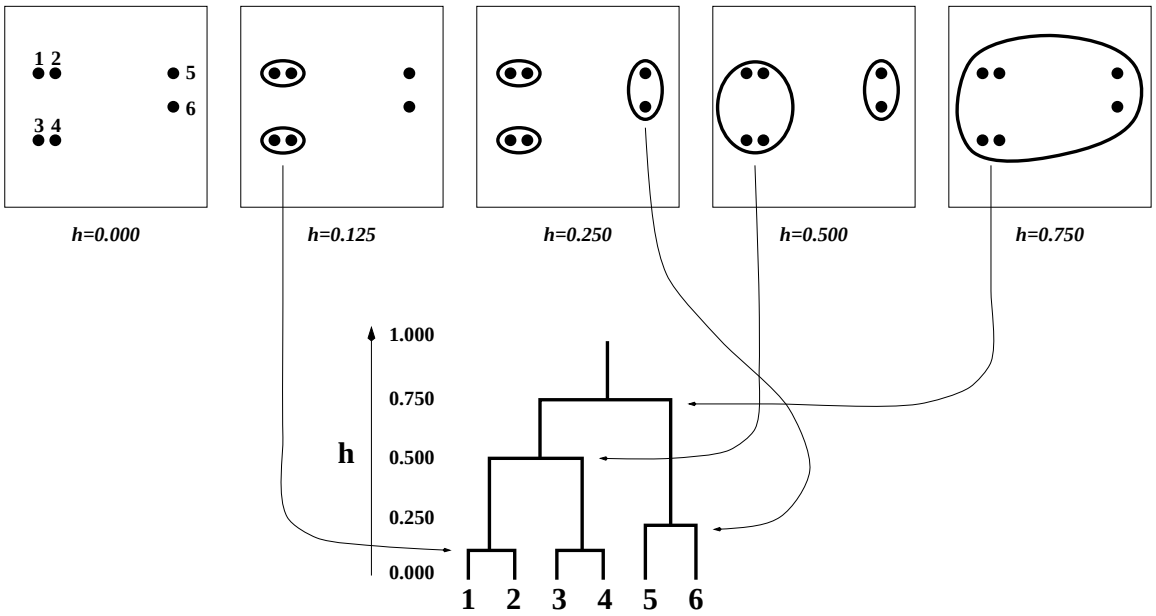

Figure 5.8: Example classification using numerical techniques. The top row shows how the system is clustered at several levels, parameterized by taxonomic level h (h is distinct from information entropy H). The classification is summarized in a taxonomic tree, or dentogram (bottom). Strong similarities between elements are indicated by grouping near the bottom of the dentogram; weaker similarities between groups are reflected in converging branches at higher levels.

Figure 5.8 provides an example of the numerical taxonomic approach. Six elements (they could be organisms, species or robots) are distributed about a two-dimensional space. The location of each element in the space is determined by the value of each trait (e.g. tail length, weight, etc.) used in the classification. Each trait corresponds to a dimension in the classification space. The goal is to build a taxonomic tree that reflects the spatial distribution of elements in the system: closely related elements should be grouped together at the bottom; similarities between groups are indicated as the branches converge at higher levels. These relations are expressed graphically in a *dentogram* (Figure 5.8, bottom).

Techniques from numerical taxonomy address the problem of how to group organisms, or groups of organisms, at various levels. At the lowest level in biological classification for instance, humans and gorillas are more likely to be grouped together

than, say, humans and dogs. But at a higher level, primates are in fact grouped together with canines in the class *mammalia*. Dentograms provide an orderly hierarchical view of the these groupings. While dentograms *per se* are not necessary for the evaluation of diversity, they are useful visualization tools and their construction provides clues for the evaluation of overall societal diversity.

Dentograms are constructed using a clustering algorithm parameterized by h, the maximum difference allowed between elements in the same group (more detail on clustering algorithms is presented later). In most applications the difference metric is normalized so that taxonomic distance between any two elements varies between 0 and 1. When $h = 1$ all elements are grouped together in one cluster (see the cluster at the top right in Figure 5.8 for example). As h is reduced from 1 down to 0 cluster boundaries change; the number of groups increases as they split into smaller clusters. The splits are reflected as branches in the dentogram. Finally, when $h = 0$ each element is a separate cluster; a "leaf" at the bottom of the dentogram "tree."

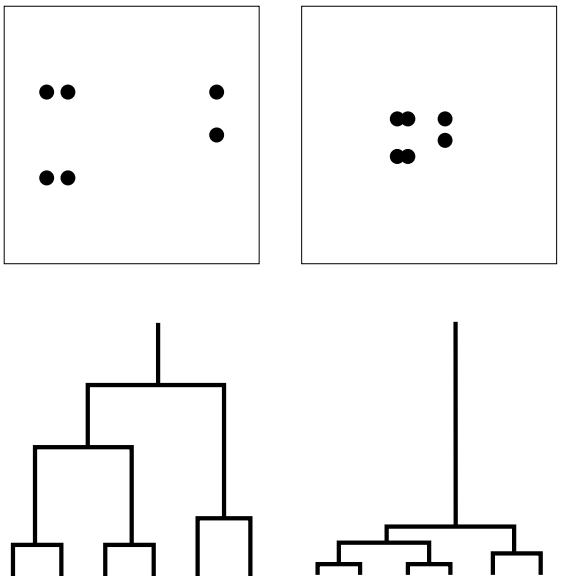

Figure 5.9: The branching structure of the dentograms for these two societies is the same. However, the more compact distribution of elements in the system on the upper right is reflected in the branches being compressed towards the bottom of the corresponding dentogram (lower right).

Dentograms can reveal subtle differences in societal structure. Figure 5.9 for example, shows two societies with the same relative arrangement of elements, but one grouping is compact while the other is spread out over a larger area. The

difference in scale is reflected in a compressed dentogram for the spatially compact society (Figure 5.9 right). Can these differences be accounted for in the evaluation of diversity?

Before addressing this, it is necessary to examine some of the details of clustering algorithms used to build a taxonomic tree. After that, the discussion returns to how these techniques can be used in the evaluation of diversity.

5.2.2 Clustering algorithms

Literally hundreds of clustering algorithms have been developed by researchers in a wide range of fields (Sneath and Sokal present a comprehensive taxonomy of clustering methods in [SS73]). One reason for the proliferation of techniques is the lack of generally agreed upon optimality criteria for evaluating the various methods. Jardine, for instance, suggests information-based metrics for biological clustering applications, but this may not be appropriate in all domains [JS71]. Because we are interested in the advantages of taxonomic representations of societal structure, the field of numerical taxonomy is an appropriate source of techniques for this research.

Most clustering methods used in numerical taxonomy are *hierarchic* as defined by Sneath and Sokal [SS73]:

> In *hierarchic* classifications any member of a lower ranking taxon is also a member of a higher ranked taxon, although not all its associates from the lower ranking taxon will necessarily be included in the higher ranking taxon. *Non-hierarchic* classifications do not exhibit ranks in which subsidiary taxa become members of larger more inclusive taxa. For traditional biological taxonomy, hierarchic classifications are required.

Another important distinction between clustering algorithms is whether or not overlap is allowed between clusters. In a *nonoverlapping* method, taxa at any one rank are mutually exclusive; a member of one group cannot also be a member of a second group at the same rank. Nonoverlapping classifications must sometimes arbitrarily assign elements to one or another equally distant group. By relaxing this constraint, *overlapping* methods allow membership in more than one taxa. The two approaches are examined below.

Nonoverlapping hierarchic clustering

Most clustering algorithms begin with a two-dimensional difference matrix. Each location in the matrix records the difference between two agents in the system to be classified. For instance, cell ij contains the value of $D(r_i, r_j)$, the difference between agents r_i and r_j (note: behavioral difference for robots is defined in the next section). Figure 5.10 illustrates an example system; the associated difference matrix is given in Table 5.2.2.

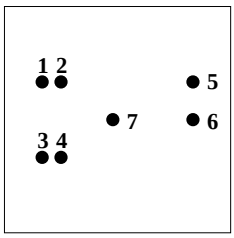

element	x	y
1	-0.35	0.25
2	-0.25	0.25
3	-0.35	-0.25
4	-0.25	-0.25
5	0.35	0.25
6	0.35	0.00
7	0.00	0.00

Figure 5.10: Example society of elements distributed in a two-dimensional space.

Table 5.1: Difference matrix for the example society. In this case, differences are calculated as Euclidean distances in the two-dimensional space.

	1	**2**	**3**	4	**5**	**6**	**7**
1	0.000	-	-	-	-	-	-
2	0.100	0.000	-	-	-	-	-
3	0.500	0.510	0.000	-	-	-	-
4	0.510	0.500	0.100	0.000	-	-	-
5	0.700	0.600	0.860	0.781	0.000	-	-
6	0.743	0.650	0.743	0.650	0.250	0.000	-
7	0.430	0.354	0.430	0.354	0.430	0.350	0.000

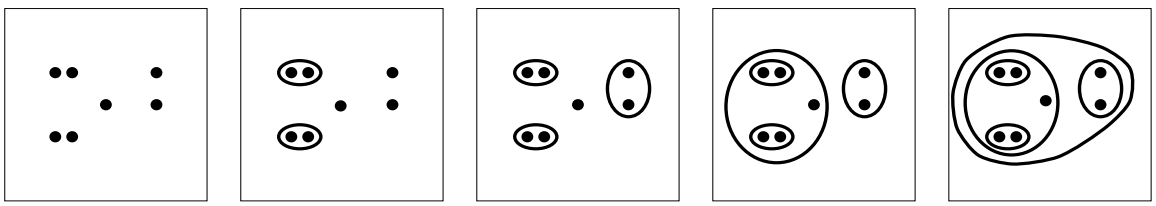

Figure 5.11: Example of hierarchic nonoverlapping clustering. Clusters are enclosed in black lines.

To compute the clusters at level h, nonoverlapping methods generally proceed in the following steps:

1. While the smallest value in the difference table, h', is less than h:
   (a) Group elements with difference h' together in a new cluster(s).
   (b) Remove elements in the new cluster(s) from the difference table.
   (c) Calculate the centroid of the new cluster(s).
   (d) Compute a new difference matrix including the new centroid(s).

The result is a hierarchy of nested clusters (Figure 5.11). Various algorithms differ in specific details, particularly in how they handle the situation where multiple pairs of elements share the same minimum difference value.

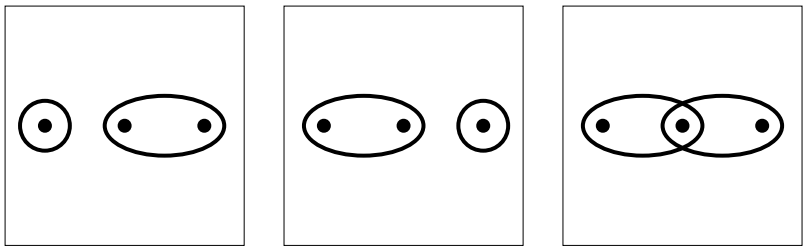

Figure 5.12: A dilemma for nonoverlapping algorithms. With which other element should the center element be grouped? Nonoverlapping algorithms must make an arbitrary choice between clustering **a** or **b**. If overlapping is allowed however, a single unambiguous solution emerges (**c**).

A key problem with nonoverlapping methods is that they must sometimes group elements arbitrarily. Figure 5.12 illustrates a case where two equally appropriate clusterings are possible (cases **a** and **b**). The nonoverlapping requirement forces an arbitrary choice of one over the other. When overlapping clusters are allowed, however, a single unambiguous solution is possible (case **c**).

Overlapping hierarchic clustering

In contrast to the nonoverlapping methods, overlapping clustering techniques allow individual elements to join more than one cluster. Overlapping methods are typically characterized by the degree of overlap allowed. Overlap can be quantified as the diameter of overlap or as the number of elements in the overlapping region.

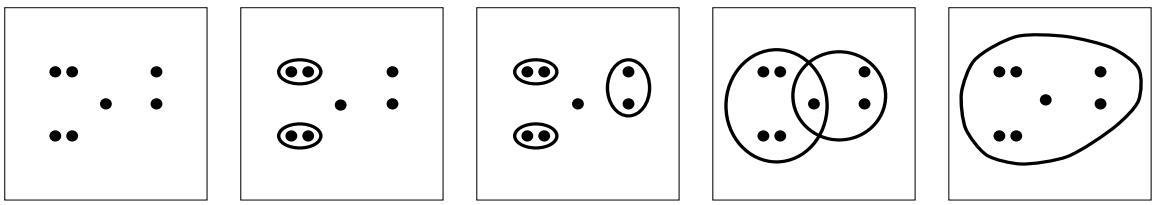

Figure 5.13: Example of hierarchic overlapping clustering. Clusters are enclosed in black lines.

The C_u clustering method is used in this research [JS71]. C_u or *u-diametric clustering* methods permit the diameter of overlap between clusters at level h to be at most uh (in this work $u = 1$). A cluster at level h is a maximally linked set such that for all elements r_i and r_j in the cluster $D(r_i, r_j) \leq h$.

Reviewing the notation presented earlier, the society of N elements to be clustered is $\mathcal{R} = \{r_1, r_2, r_3, \ldots r_N\}$. The society will be divided into M possibly overlapping clusters $\mathcal{C} = \{c_1, c_2, c_3, \ldots c_M\}$. The C_u algorithm for clustering at level h proceeds in the following sequence:

1. Initialize N clusters with $c_i = \{r_i\}$.
2. For each cluster c_i:
   (a) For each r_j (except r_i) in $\mathcal{R}$:
      i. If $(D(r_j, r_k) \leq h)$ for every r_k already in c_i) add element r_j to cluster c_i.
3. Discard redundant clusters.

An example society grouped using C_u clustering is presented in Figure 5.13. The clusterings for several values of h are illustrated with h increasing from left to right. Notice in the fourth diagram that the element in the middle of the space is claimed by two clusters (contrast this with Figure 5.11). This clustering technique is *hierarchic* because elements grouped at one level, or value of h, are also members of higher level taxa. In addition, the taxa (clusters) become larger and more inclusive at higher levels.

The AutoClass clustering program

In addition to the techniques outlined above, there are a number of other clustering algorithms in current use. One of these, AutoClass, was evaluated for use in this research [CS96]. AutoClass is a Bayesian nonhierarchic nonoverlapping clustering

algorithm available in the public domain. According to Cheeseman, one of AutoClass' developers, "the goal of Bayesian unsupervised classification is to find the most probable set of class descriptions given the data and prior expectations."

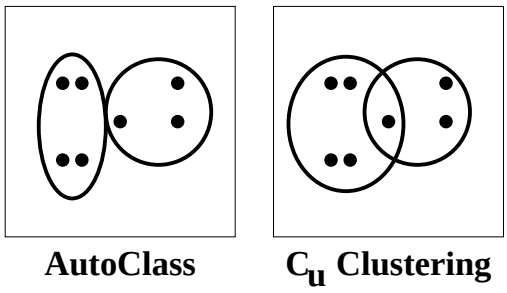

Figure 5.14: A comparison of clusterings generated by AutoClass (left) and C_u clustering (right) for the example data. $h = 0.51$ for C_u clustering.

The AutoClass system was used to classify several datasets for comparison with C_u clustering. The AutoClass classification for the system given in Figure 5.10 is provided in Figure 5.14. The result is compared with a C_u clustering of the same system. In these examples h for C_u clustering is set to 0.51; the clusterings indicate element similarities at a particular scale. Note that h can be varied, however, to examine the structure of the society at any scale.

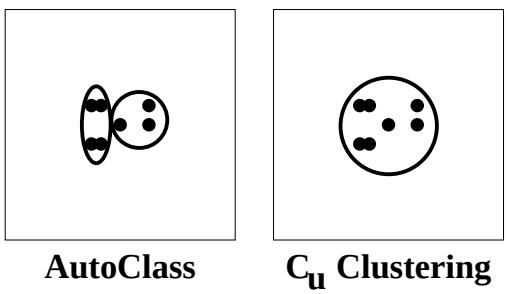

Figure 5.15: A set of clusterings for a more compact society. The locations of elements in this example are scaled by $\frac{1}{2}$ in comparison to the previous example. $h = 0.51$ for C_u clustering.

In another evaluation, AutoClass was used to classify a more spatially compact system (for this evaluation, the previous data values are multiplied by $\frac{1}{2}$). The same axes were used in both evaluations, only the locations of the elements in the space were changed. Results of the classification are given in Figure 5.15, and for comparison the same system is classified using the C_u algorithm. The groups found

by AutoClass are the same as in the previous example. In contrast, the C_u algorithm groups all of the elements together in one cluster.

Two additional tests were conducted to evaluate AutoClass and C_u clustering at extreme ends of the spectrum. In the first test the data values used in Figure 5.14 were multiplied by 100000 and the resulting system was again classified by AutoClass and C_u clustering. As before, AutoClass groups the elements in the same way. C_u clustering, with $h = 0.51$ places each element in a separate group. Finally, the data values were multiplied by $\frac{1}{100000}$ and grouped by AutoClass and C_u clustering. Again, AutoClass finds the same groups. In contrast C_u clustering places all elements in the same group.

Discussion on clustering algorithms

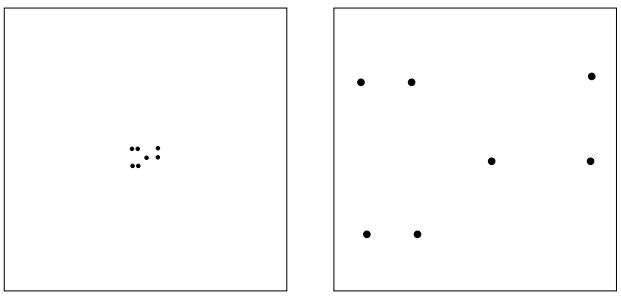

Figure 5.16: n example of two societies differing only in the degree of difference between elements. Axes used in both systems are the same. The relative position of each element is the same in each society, but in the system on the right they are much further apart. A useful diversity metric should distinguish between these two societies.

It is impossible to distinguish between the systems illustrated in Figure 5.16 on the basis of clusterings generated by AutoClass because AutoClass is not sensitive to the degree of difference between elements. It would cluster both systems identically. In fairness to AutoClass, it is likely that its designers sought to avoid sensitivity to the spatial scope of the data under evaluation. This feature is desirable in classification applications involving a single dataset. This work, however, is concerned with *comparing* several datasets that may vary in spatial extent (e.g. the systems in Figure 5.16).

The spatial extent of elements in a taxonomic space is a reflection of the degree of difference between agents. It has already been pointed out (in Section 5.1.3) that such differences are important in the evaluation of diversity, especially for distinguishing

between societies with similar structure and numbers of elements but with differing spatial size. Overall spatial size for instance, is the only difference between the societies illustrated in Figure 5.16. Because differentiating between societies on the basis of their extent is important in the evaluation of diversity, this work utilizes C_u clustering in diversity calculation.

Note that sensitivity to the degree of difference between elements in hierarchic clustering depends on h. Because h is a *parameter* of the clustering algorithm, it can be varied to examine clusterings at any scale. Hierarchic algorithms are, in effect, variable power clustering microscopes. For values of h near zero the tiniest difference between elements will cause them to be grouped separately, while the clusterings at large values of h reveal societal structure at a macroscopic level. This feature is exploited in the development of a diversity measure sensitive to differences in the spatial size of societies.

5.2.3 Hierarchic social entropy

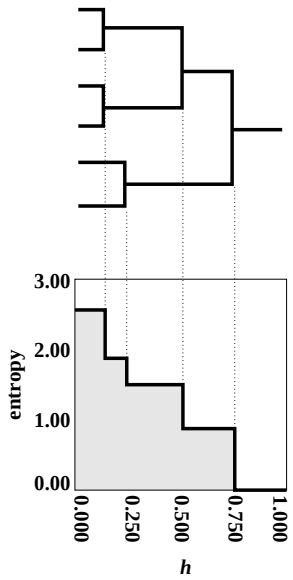

Figure 5.17: Entropy depends on h. Note that changes in entropy correspond to the branch points in the dentogram. (For easier reference, the dentogram is rotated 90 degrees.)

Now consider how tools from numerical taxonomy can be applied to the measurement of diversity. The discussion of hierarchic clustering algorithms above described how the number and size of clusters depend on h. But how is social entropy impacted

by changes in h? Since the partitioning of a society is based on h the entropy also depends on it. An example of the relationship is illustrated in Figure 5.17. Entropy changes in discrete steps as h increases. Note that points where change occurs correspond to branch points in the dentogram.

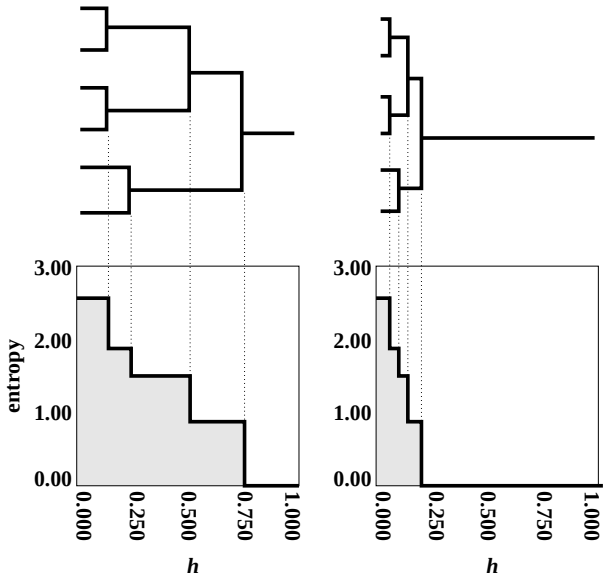

Figure 5.18: A comparison of entropy versus h for for two societies.

Compare the dentograms and entropy plots of the two societies in Figure 5.18. As in the earlier example, the two groups have the same relative structure, but the society represented on the right is more compact, resulting in branching compressed towards the bottom of the tree. The difference in scale is also readily apparent in the plots of entropy. Entropy drops to zero much more quickly in the plot corresponding to the compact society. Because the value of simple entropy depends significantly on h when hierarchic clustering is used, we augment the notation to account for this:

$$H(\mathcal{R}, h) \quad = \quad H(\mathcal{R}) \text{ for clustering at taxonomic level } h \tag{5.3}$$

H is a function of $\mathcal{R}$ *and* h because the classification of agents into groups, and therefore the entropy, depends on them both. This highlights the fact that **the entropy of a particular clustering is only a snapshot of the society's diversity**. A comprehensive evaluation of diversity should account for clustering at all taxonomic levels. This is easily accomplished using the area under the entropy plot

as a measure of diversity. This augmented metric, called *hierarchic social entropy*, is defined as:

$$S(\mathcal{R}) = \int_0^\infty H(\mathcal{R}, h)dh \tag{5.4}$$

where $\mathcal{R}$ is the robot society under evaluation, h is a parameter of the clustering algorithm indicating the maximum difference between any two agents in the same group and $H(\mathcal{R}, h)$ is the simple entropy of the society for the clustering at level h. Note that as $h \to \infty$ a point is reached where all elements are clustered in the same group (the maximum taxonomic distance). $H(\mathcal{R}, h)$ drops to 0 at this point. In the behavioral difference measure used in this work, the maximum possible difference between elements is fixed at 1.0, so the upper limit of the integration is 1 rather than ∞ as in the general case.

To clarify how hierarchic entropy is computed, several examples are provided in the next section.

5.2.4 Why hierarchic social entropy is a useful metric

Hierarchic social entropy is a continuous ratio measure; it has an absolute zero (when all elements are identical) and equal units (bits). This enables a total ordering of societies on the basis of diversity. It also provides for quantitative results of the form "$\mathcal{R}_b$ is *twice* as diverse as $\mathcal{R}_a$." This is a significant advantage over the categorization of systems as simply "homogeneous" or "heterogeneous." Several other useful properties of hierarchic entropy are examined below.

Hierarchic social entropy can distinguish differences between societies regardless of scale. Societies with infinitesimally small differences are compared as easily and precisely as systems spanning millions of units. This property is demonstrated with an example. Figure 5.19 illustrates two societies of three elements arranged in triangles. In both cases the two elements on the left are spaced a distance x apart. A third element is placed either $2x$ or $4x$ from the other elements in societies $\mathcal{R}_{2x}$ and $\mathcal{R}_{4x}$ respectively. Because hierarchic social entropy is scale invariant, it is able to distinguish between these two systems for all values of x. This will be demonstrated for $x = \frac{1}{1000000}$, for $x = 1000000$ and proven for all x.

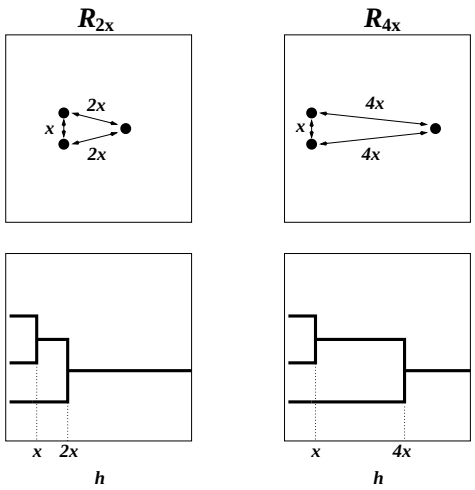

Figure 5.19: These two example systems are used to demonstrate how hierarchic social entropy can distinguish differences between societies regardless of scale. Spacing between the elements is parameterized by x (top). Because hierarchic entropy is scale invariant it can distinguish between the two societies regardless of the value of x. Dentograms (bottom) illustrate the values of h where clusterings change.

First, observe that due to the spacing of the elements, there are three distinct C_u clusterings for each system (depending on h).[1] For society $\mathcal{R}_{2x}$, the three elements are placed in three separate clusters when $0 \leq h < x$. Two clusters are present when $x \leq h < 2x$. Finally, all three elements are grouped together in one cluster when $2x \leq h$. The groupings over all three ranges are illustrated in Figure 5.20 (groupings

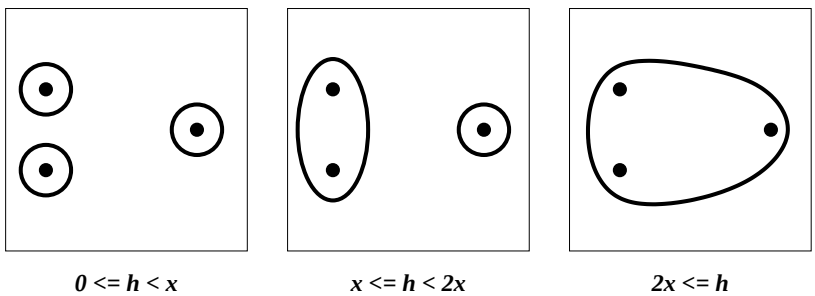

Figure 5.20: C_u clusterings of $\mathcal{R}_{2x}$ for different values of h.

are similar for $\mathcal{R}_{4x}$ except the final clustering does not occur until $h \geq 4x$). The simple entropy for each clustering of $\mathcal{R}_{2x}$ is

[1]Note: the AutoClass program groups all elements together in the same cluster for both societies when $x = \frac{1}{1000000}$ and when $x = 1000000$. A diversity metric based on AutoClass clusterings would not distinguish between these two societies.

$$
\begin{aligned}
H(\frac{1}{3},\frac{1}{3},\frac{1}{3}) &= 1.585 \text{ for } 0 \leq h < x \\
H(\frac{2}{3},\frac{1}{3}) &= 0.811 \text{ for } x \leq h < 2x \\
H(\frac{3}{3}) &= 0.000 \text{ for } 2x \leq h
\end{aligned}
$$

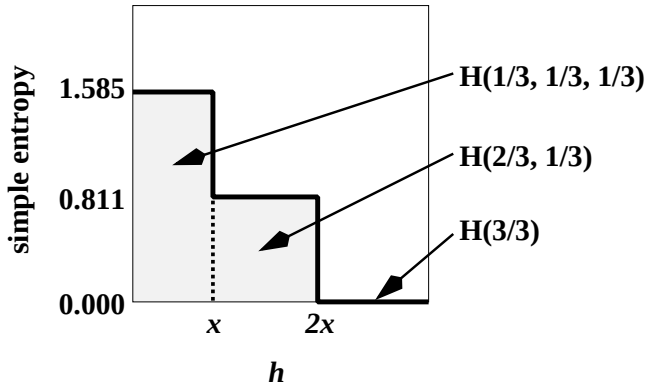

Figure 5.21: Simple entropy of $\mathcal{R}_{2x}$ as a function of h. There are three distinct regions with different values.

These values and the regions over which they apply are illustrated in Figure 5.21. Similarly, the simple entropy for each clustering of $\mathcal{R}_{4x}$ is

$$
\begin{aligned}
H(\frac{1}{3},\frac{1}{3},\frac{1}{3}) &= 1.585 \text{ for } 0 \leq h < x \\
H(\frac{2}{3},\frac{1}{3}) &= 0.811 \text{ for } x \leq h < 4x \\
H(\frac{3}{3}) &= 0.000 \text{ for } 4x \leq h
\end{aligned}
$$

Now, suppose $x = \frac{1}{1000000}$. Can hierarchic entropy distinguish between these two systems? First we calculate the hierarchic entropy of society $\mathcal{R}_{2x}$. Recall the definition of hierarchic social entropy (Equation 5.4):

$$
S(\mathcal{R}) = \int_0^\infty H(\mathcal{R},h)dh
$$

As was pointed out above, $H(\mathcal{R}, h)$ takes on distinct values over three regions depending on h. Therefore, the integral can be broken into parts corresponding to these regions:

$$\int_0^\infty H(\mathcal{R}_{2x}, h)dh = \int_0^x H(\mathcal{R}_{2x}, h)dh + \int_x^{2x} H(\mathcal{R}_{2x}, h)dh + \int_{2x}^\infty H(\mathcal{R}_{2x}, h)dh$$

Substituting $\frac{1}{1000000}$ for x and the simple entropy values above for $H(\mathcal{R}_{2x}, h)$, we have

$$\begin{aligned}
\int_0^\infty H(\mathcal{R}_{2x}, h)dh &= \int_0^{\frac{1}{1000000}} H(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}) + \int_{\frac{1}{1000000}}^{\frac{2}{1000000}} H(\frac{2}{3}, \frac{1}{3}) + \int_{\frac{2}{1000000}}^\infty H(\frac{3}{3}) \\
&= \int_0^{\frac{1}{1000000}} 1.585 + \int_{\frac{1}{1000000}}^{\frac{2}{1000000}} 0.811 + \int_{\frac{2}{1000000}}^\infty 0.000 \\
&= 2.396 \times 10^{-6}
\end{aligned}$$

The hierarchic social entropy of $\mathcal{R}_{2x}$ is 2.396×10^{-6}. The calculation for $\mathcal{R}_{4x}$ is similar:

$$\begin{aligned}
\int_0^\infty H(\mathcal{R}_{4x}, h)dh &= \int_0^x H(\mathcal{R}_{4x}, h)dh + \int_x^{4x} H(\mathcal{R}_{4x}, h)dh + \int_{4x}^\infty H(\mathcal{R}_{4x}, h)dh \\
&= \int_0^{\frac{1}{1000000}} H(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}) + \int_{\frac{1}{1000000}}^{\frac{4}{1000000}} H(\frac{2}{3}, \frac{1}{3}) + \int_{\frac{4}{1000000}}^\infty H(\frac{3}{3}) \\
&= \int_0^{\frac{1}{1000000}} 1.585 + \int_{\frac{1}{1000000}}^{\frac{4}{1000000}} 0.811 + \int_{\frac{4}{1000000}}^\infty 0.000 \\
&= 4.018 \times 10^{-6}
\end{aligned}$$

For system $\mathcal{R}_{4x}$ we have $S(\mathcal{R}_{4x}) = 4.018 \times 10^{-6}$. Therefore when $x = \frac{1}{1000000}$ $S(\mathcal{R}_{2x}) < S(\mathcal{R}_{4x})$ and $\mathcal{R}_{4x}$ is 1.68 times more diverse than $\mathcal{R}_{2x}$.

What if $x = 1000000$? For $\mathcal{R}_{2x}$ the computation proceeds as follows:

$$\begin{aligned}
\int_0^\infty H(\mathcal{R}_{2x}, h)dh &= \int_0^x H(\mathcal{R}_{2x}, h)dh + \int_x^{2x} H(\mathcal{R}_{2x}, h)dh + \int_{2x}^\infty H(\mathcal{R}_{2x}, h)dh \\
&= \int_0^{1000000} H(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}) + \int_{1000000}^{2000000} H(\frac{2}{3}, \frac{1}{3}) + \int_{2000000}^\infty H(\frac{3}{3})
\end{aligned}$$

$$
\begin{aligned}
&= \int_0^{1000000} 1.585 + \int_{1000000}^{2000000} 0.811 + \int_{2000000}^{\infty} 0.000 \\
&= 2.396 \times 10^6
\end{aligned}
$$

Similarly, the hierarchic entropy for system $\mathcal{R}_{4x}$ is 4.018×10^6. So when $x = 1000000$ we again have $S(\mathcal{R}_{2x}) < S(\mathcal{R}_{4x})$; society $\mathcal{R}_{4x}$ is again 1.68 times more diverse than $\mathcal{R}_{2x}$

In fact, $S(\mathcal{R}_{2x}) < S(\mathcal{R}_{4x})$ holds for all values of $x > 0$:

$$
\begin{aligned}
\int_0^{\infty} H(\mathcal{R}_{2x}, h)dh &\stackrel{?}{<} \int_0^{\infty} H(\mathcal{R}_{4x}, h)dh \\
\int_0^{x} H(\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{1}{3}) + \int_x^{2x} H(\tfrac{2}{3}, \tfrac{1}{3}) + \int_{2x}^{\infty} H(\tfrac{3}{3}) &\stackrel{?}{<} \int_0^{x} H(\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{1}{3}) + \int_x^{4x} H(\tfrac{2}{3}, \tfrac{1}{3}) + \int_{4x}^{\infty} H(\tfrac{3}{3}) \\
\int_0^{x} 1.585 + \int_x^{2x} 0.811 + \int_{2x}^{\infty} 0.000 &\stackrel{?}{<} \int_0^{x} 1.585 + \int_x^{4x} 0.811 + \int_{4x}^{\infty} 0.000 \\
1.585x + 0.811x + 0.000 &\stackrel{?}{<} 1.585x + 0.811 \times 3x + 0.000 \\
2.396x &< 4.018x
\end{aligned}
$$

In addition to scale invariance, hierarchic entropy benefits from several other advantages. **Hierarchic entropy addresses a key weakness of simple social entropy by accounting for continuous differences between elements in the society.** Figure 5.22 illustrates the kind of difference in societal structure hierarchic entropy can distinguish. In an earlier example, simple social entropy could not resolve differences between these systems (Figure 5.22). However, when hierarchic social entropy is employed, the measured diversity of the three systems increases linearly as the one agent is positioned further and further away. As one would expect, the difference in diversity between systems 5.22a and 5.22b is much smaller than that between 5.22b and 5.22c. This is not necessarily the case when simple entropy is used (as the earlier example illustrates)

Hierarchic entropy preserves the basic properties of simple social entropy. Hierarchic entropy is a more general metric than simple entropy, subsuming the properties of H at each taxonomic level h. In the case where difference between agents is binary, (either alike or unlike), Equation 5.4 degenerates to $H(\mathcal{R})$ (simple entropy) because the clustering does not depend on h. However, when continuous differences are important, hierarchic entropy can resolve structural difference in so-

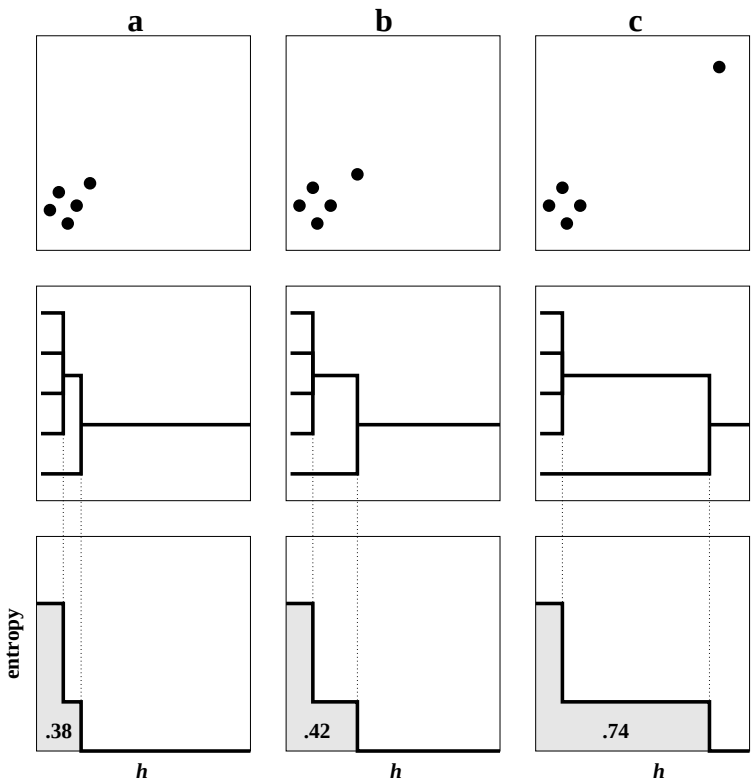

Figure 5.22: Hierarchic social entropy (bottom) is computed for three societies (top). The calculated value increases as the element on the upper right is positioned further away from the group. Dentograms for the groups are also displayed (middle row).

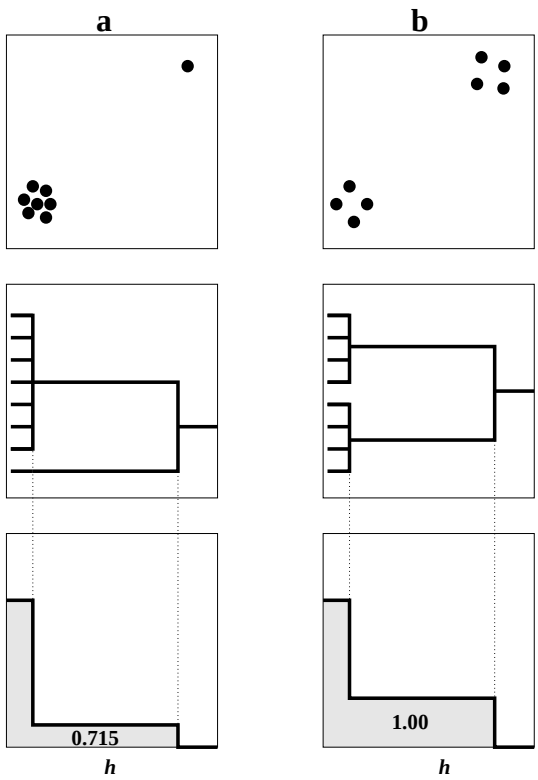

Figure 5.23: Hierarchic social entropy retains the basic properties of simple entropy. The computed value (bottom) depends on the distribution of elements in the groups. Dentograms for the two groups are also displayed (middle row).

cieties that simple social entropy cannot.

Figure 5.23 shows how the basic properties of social entropy are preserved with hierarchic entropy. In this example two groups are located a fixed distance apart in the classification space. The two societies pictured differ only in the distribution of elements between the groups. Hierarchic entropy properly captures the increased diversity of the system with agents distributed equally between the groups.

5.3 Behavioral difference

To summarize the chapter so far, hierarchic clustering is a means of dividing a society into groups of behaviorally equivalent agents at a particular taxonomic level. Diversity is evaluated at each taxonomic level based on the number of groups and the number of robots in each group at that level. Integrating the diversity across all taxonomic levels produces an overall measure of diversity for the system. Previous sections have described the overall diversity metric and algorithms for clustering the agents into groups. This section focuses on the difference metric used for clustering.

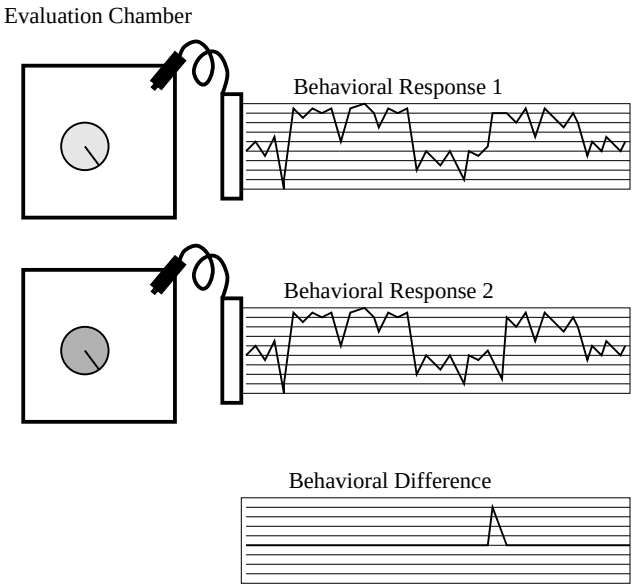

Figure 5.24: [Evaluating behavioral difference using an idealized "evaluation chamber." Robots are evaluated in the chamber (left), where their response to every situation is recorded as a trace (readout, right). The behavioral difference between two agents is the difference between their traces (bottom). A single quantitative value is given by integrating the difference.

How should the behavior of two agents be compared? One possibility would be to

evaluate their difference in an "evaluation chamber" in which the robots are exposed to all situations and their their responses recorded. Even though it is unlikely such a chamber could be built, the analogy is useful. Figure 5.24 illustrates the procedure. As the agents are exposed to various situations, responses are recorded as a trace. After the experiment is concluded, the traces are compared to evaluate the difference between agents. In the figure, the horizontal axis of the traces represents all distinct perceptual situations a robot might experience, while the vertical component encodes the agent's response.

Since a real evaluation chamber would be practically impossible to build, an alternative method for evaluating behavioral difference is proposed. The technique advocated here is to look for differences in the agents' behavioral coding. In many cases (e.g. [BBC+95, Mat92, GM97]) robot behavior is coded statically ahead of time, thus individuals may be directly compared by evaluating their behavioral configuration. Learning multi-robot systems (e.g. [Bal97c, Mat94]) pose a challenge because their behavior evolves over time. To avoid that problem in this research, the policies of learning agents are evaluated after agents converge to stable behavior.

This approach depends on three key assumptions:

Assumption 1: At the time of comparison, the robots' policies are fixed and deterministic.

Assumption 2: The robots under evaluation are substantially mechanically similar: differences in overt behavior are influenced more significantly by differences in policy than by differences in hardware.

Assumption 3: Differences in policy are correlated with differences in overt behavior.

If these conditions are not met in a particular multi-robot system, the approach may not be appropriate. But the assumptions are reasonable for the conditions of this research, namely: experiments conducted on mechanically similar robots built on the same assembly line. Additionally, the robots all run identical control system software. The control systems running on the robots differ only in the data specifying each agent's policy. The comparison of these policies is the crux of the approach.

5.3.1 Example: multi-robot foraging

The objective is to show how behavioral difference can be evaluated by examining differences in robots' behavioral coding. Before proceeding, an example encoding is

Figure 5.25: An example task for a multiagent robot team. The robots are to collect red and blue objects and place them into colored bins. The task is similar to the 1997 AAAI Mobile Robot Contest's "Find Life on Mars" event. The object in the foreground is a "black rock" obstacle.

presented. Although this example describes a particular robot architecture, this is only for illustration, the method is applicable to other architectures as well.

Consider how behaviors could be designed for a team of foraging robots (Figure 5.25). The task is to collect colored objects (red) and place them into colored bins (red and blue). For this example, one agent will be programmed to place the objects in the red bin, while the other will deposit them in the blue bin.[2]

In this approach to behavioral configuration, the agent is provided several behavioral assemblages that correspond to steps in achieving the task (e.g. *wander*, *acquire*, *deliver*, and so on). Binary perceptual cues are used to sequence the robot through the steps in achieving the task.

The agents are provided with the perceptual features enumerated in Table 5.2. At the behavior selection level, the robot's perception can be represented by four bits (one bit per perceptual feature). Given the perceptual state, the robot selects from one of the four behaviors listed in Table 5.3. Decomposing the task into a state/action space enables a robot's policy to be enumerated by pairing perceptual states with actions. Some of the 16 states are never actually encountered since it is impossible for an agent to be simultaneously in the red and blue delivery zones.

[2]This task is a simplified version of the task for robots in the AAAI-97 contest. The simplification is necessary in order to allow a complete enumeration of the robots' policies.

Table 5.2: Perceptual features available to the foraging robots. Each feature is equivalent to one bit of information; the entire perceptual state is a four-bit value.

perceptual feature	meaning
`red_visible`	a red attractor is visible.
`red_in_gripper`	a red attractor is in the gripper.
`close_to_red_bin`	close enough to the red delivery area to drop an attractor in it.
`close_to_blue_bin`	close enough to the blue delivery area to drop an attractor in it.

Table 5.3: Behaviors the robots select from in accomplishing the foraging task.

behavior	robot activity
wander	Search the environment for attractors.
acquire_red	Proceed to the closest red object and grasp it.
deliver_blue	Go the the blue delivery area.
deliver_red	Go to the red delivery area.

Using this approach, policies for the two robots are now described. One of the robots collects red objects and places then in the red bin, while the other places them in the blue bin. The policy for robot r_{red} is to search for red attractors using the *wander* behavior. When it sees an attractor, it activates the *acquire_red* behavior. Once it has grasped the object, it uses the *deliver_red* behavior to go to the red bin. Robot r_{blue} is similar, except it delivers to the blue bin instead. Policies for the two agents are enumerated in Table 7.2. The behaviors in the center of the table are activated when the corresponding perceptual situations on the left are encountered.[3]

The actions selected by the agents described above differ in six of the states. In the case where the robots have a red object in their gripper but aren't close to a bin, they choose different actions (to either go to the red or blue bin). When they are close to the correct bin, they both drop the attractor and resume the *wander* behavior. The next section explains how a numerical value can be assigned this behavioral difference.

[3]Note that in the *wander* behavior, the robot's gripper opens automatically. A transition to the *wander* behavior causes the robot to drop the attractor and begin a new search.

Table 5.4: The policies of two foraging robots. Robot r_{red} collects red objects and places them in the red bin, Robot r_{blue} collects red objects and places then in the blue bin. Differences between the actions are listed on the right. The state bits represent, from left to right, `red_visible`, `red_in_gripper`, `close_to_red_bin`, and `close_to_blue_bin`. Impossible states are indicated with an asterisk (*).

state	robot r_{red} action	robot r_{blue} action	response difference
0000	*wander*	*wander*	0.0
0001	*wander*	*wander*	0.0
0010	*wander*	*wander*	0.0
0011*	*wander*	*wander*	0.0
0100	*deliver_red*	*deliver_blue*	1.0
0101	*deliver_red*	*wander*	1.0
0110	*wander*	*deliver_blue*	1.0
0111*	*wander*	*wander*	0.0
1000	*acquire_red*	*acquire_red*	0.0
1001	*acquire_red*	*acquire_red*	0.0
1010	*acquire_red*	*acquire_red*	0.0
1011*	*acquire_red*	*acquire_red*	0.0
1100	*deliver_red*	*deliver_blue*	1.0
1101	*deliver_red*	*wander*	1.0
1110	*wander*	*deliver_blue*	1.0
1111*	*wander*	*wander*	0.0

5.3.2 Definition of behavioral difference

To facilitate the discussion, the following additional symbols and terms are defined:

- i_j is r_j's **perceptual state**.
- a_j is the action (behavioral assemblage) selected by r_j's control system based on the input i_j.
- π_j is r_j's policy; $a_j = \pi(i_j)$.
- p_j^i is the number of times r_j has encountered perceptual state i divided by the total number of times all states have been encountered. Experimentally, p_j^i is computed *post facto*.

The approach is to evaluate behavioral difference by comparing the robots' policies. The two foraging robots introduced earlier, for example, exhibit behavioral differences that are reflected in and caused by their differing policies. In the terminology introduced above, i represents the perceptual features an agent uses to selectively activate behaviors. In the case of the foraging robots, assign a bit to each perceptual feature, so, for example, $i =$ `0001` indicates that only the last perception

(`close_to_blue_bin`) is activated. For the foraging robots a is the activated behavior (e.g. *wander*, *deliver*).

Definition 1: r_a and r_b, are **absolutely behaviorally equivalent** iff they select the same behavior in every perceptual state.

In complex systems with perhaps thousands of states and hundreds of actions it may also be useful provide a scale of equivalence. This would allow substantially similar agents to be grouped in the same cluster even though they differ by a small amount. The approach is to compare two robots, r_a and r_b, by integrating the differences between their responses, $\mid \pi_a(i) - \pi_b(i) \mid$ over all perceptual states i. If the action is a single-dimension scalar, as in a motor current for instance, the difference can be taken directly. However, complex actions like *wander* and *acquire* are treated as nominal values with response difference defined as 0 when $\pi_a(i) = \pi_b(i)$ and 1 otherwise. This approach is often used in classification applications to quantify difference between nominal variables (e.g. eye color, presence or absence of a tail, etc.). Using this notation, a simple behavioral difference metric can be defined as:

$$D'(r_a, r_b) = \frac{1}{n} \int \mid \pi_a(i) - \pi_b(i) \mid di \tag{5.5}$$

or for discrete state/action spaces:

$$D'(r_a, r_b) = \frac{1}{n} \sum_i \mid \pi_a(i) - \pi_b(i) \mid \tag{5.6}$$

where $\frac{1}{n}$ is a normalization factor to ensure the difference ranges from 0 to 1. In the case of the discrete sum, n corresponds to the number of possible states. If r_a and r_b select identical outputs (a) in all perceptual states (i), then $D'(r_a, r_b) = 0$. When r_a and r_b select different outputs in all cases $D'(r_a, r_b) = 1$. In the numerical taxonomy literature, this difference is called the *mean character difference* [SS73]. The calculation parallels the idealized evaluation chamber procedure introduced earlier (Figure 5.24).

Equations 5.5 and 5.6 weigh differences equally across all perceptual states. This may be problematic for agents that spend large portions of their time in a small portion of the states. Consider two foraging robots that differ only in their reaction

to blue attractors. If, in their environment, no blue attractors are present the agents would appear to an observer to have identical policies.

There may be other important reasons certain states are never visited. In learning a policy, for instance, the robots might discover in early trials that certain portions of the state space should be avoided due to large negative rewards. Because these portions of the space are avoided, the agents will not refine their policies there, but avoid them entirely. It is entirely possible for the agents to differ significantly in these portions of the space even though they may appear externally to behave identically.

To address this, the response differences in states most frequently visited should be emphasized while those that are infrequently experienced should be de-emphasized. This is accomplished by multiplying the response difference in each situation by the proportion of times that state was visited by each agent $(p_a^i + p_b^i)$. Formally, **behavioral difference** between two robots r_a and r_b is defined as:

$$D(r_a, r_b) = \int \frac{(p_a^i + p_b^i)}{2} \mid \pi_a(i) - \pi_b(i) \mid di \tag{5.7}$$

or in discrete spaces

$$D(r_a, r_b) = \sum_i \frac{(p_a^i + p_b^i)}{2} \mid \pi_a(i) - \pi_b(i) \mid \tag{5.8}$$

When r_a and r_b select differing outputs in a given situation, the difference is normalized by the joint proportion of times they have experienced that situation.

As an example of how behavioral difference is calculated, suppose the robots introduced earlier are evaluated in an experimental run.[4] During the experiment, the number of times each agent visits each state is recorded. This log, along with the response differences listed in Table 7.2 can be used to compute the behavioral difference between the two agents. The calculation is illustrated in Table 5.5. The number of times each agent visited each state is enumerated, then used to compute p^i for each robot for each state. The normalized behavioral difference at each state is listed in the right column, and summed at the lower right. The value in the lower right-hand corner, 0.55, is the behavioral difference between robots r_{red} and r_{blue}.

The measure of behavioral difference provides for the following definitions:

[4]This example experiment is for illustrative purposes only.

Table 5.5: Sample evaluation of the behavioral difference between the two agents whose policies are listed in Table 4.3. The number of times each state was visited by each agent is listed and used to compute p^i for each state for each robot. In turn, the proportion of visits to each state is used to normalize the response difference between the agents. Note: these values were not gathered from experiment, they are presented for example only.

state	times r_{red} visited	p^i_{red}	times r_{blue} visited	p^i_{blue}	difference	normalized difference
0000	100	0.1	100	0.1	0.0	0.00
0001	-	-	100	0.1	0.0	0.00
0010	100	0.1	-	-	0.0	0.00
0011	-	-	-	-	0.0	0.00
0100	200	0.2	200	0.2	1.0	0.20
0101	-	-	200	0.2	1.0	0.10
0110	200	0.2	-	-	1.0	0.10
0111	-	-	-	-	0.0	0.00
1000	100	0.1	100	0.1	0.0	0.00
1001	-	-	100	0.1	0.0	0.00
1010	100	0.1	-	-	0.0	0.00
1011	-	-	-	-	0.0	0.00
1100	100	0.1	100	0.1	1.0	0.10
1101	-	-	100	0.1	1.0	0.05
1110	100	0.1	-	-	1.0	0.00
1111	-	-	-	-	0.0	0.00
totals	1000	1.0	1000	1.0	6.0	0.55

Definition 2: r_a and r_b, are **ϵ-equivalent** iff $D(r_a, r_b) < \epsilon$.

Definition 3: $\equiv_\epsilon$ indicates ϵ-equivalence, $r_a \equiv_\epsilon r_b$ means r_a and r_b are ϵ-equivalent.

Definition 4: A robot society, $\mathcal{R}$, is **ϵ-homogeneous** iff for all $r_a, r_b \in \mathcal{R}$, $r_a \equiv_\epsilon r_b$

ϵ in these definitions is closely tied to the parameter h used in C_u clustering. A classification at taxonomic level h will consist of h-homogeneous clusters.

5.3.3 Limitations of the approach

An important limitation of this approach to evaluating behavioral difference is the requirement for an agent's policy to be represented as a deterministic function, e.g. $a = \pi(i)$. This is reasonable for the analysis of policies developed using reinforcement learning techniques since once learning is complete, the policies are fixed. But the approach does not address robots utilizing FSAs for behavioral sequencing. An FSA might generate a different output (action) in the same perceptual state, depending on the sequence of inputs up to that point. To address the problem a quantitative technique for comparing FSAs is required. One avenue of approach would be a comparative analysis of the "languages" (actually sequences of perceptual state) accepted by two agents under evaluation. This is beyond the scope of the present investigation however.

Another potential problem is the implicit assumption that two robots being compared sequence through the state space at the same rate. As an example, two robots might have the same policy but because one of them is equipped with a very slow computer it changes state much less frequently. An analysis of behavioral difference based solely on policy difference could not differentiate between these two agents. The problem could be addressed by augmenting the difference metric with an additional "computational resource" factor. The issue was not a problem in this research because each robot is provided with identical computing resources.

5.4 Discussion and summary

Researchers in multiagent systems need new tools for measuring agent difference and diversity in the artificial societies they study. Without such metrics it is impossible to evaluate the quantitative impact of task on diversity or to correlate diversity with performance. To address these issues, this chapter introduces

- a mathematical expression for the **behavioral difference** between two robots,
- rigorous definitions of **behavioral homogeneity** and **heterogeneity** in multi-robot teams, and
- **social entropy**, a new measure of robot team **behavioral diversity**.

Behavioral difference is evaluated by comparing the policies of robots and normalizing the differences with respect to the proportion of times the two agents experience each perceptual state. The behavioral difference measure can be used in turn by clustering algorithms to divide the society into groups of similar agents.

The field of numerical taxonomy provides is a rich source of methods for this work because biologists face many of the same challenges. The hierarchic clustering methods used in taxonomy are advantageous because the generated hierarchies provide more information about the structure of the society than a simple clustering. C_u clustering, a hierarchic overlapping clustering algorithm was adopted for this research.

The diversity of a robot society is computed using the classification of agents into groups. The social entropy (diversity) of a society is computed based on the number of groups and the size of each group. Use of the metric is illustrated in several example evaluations.

The value of social entropy as a diversity metric may be limited when difference between agents is continuous. Simple entropy does not account for the degree of difference between agents, or the structure of the society in the classification space. To address this, a more comprehensive measure, hierarchic social entropy is introduced. In short, hierarchic social entropy is the average entropy of a society across all taxonomic levels ($0 \leq h \leq 1$).

Both simple and hierarchic social entropy are used in the experiments reported in later chapters of this work. Simple entropy is employed when obvious distinctions between agents can be made or when the behavioral difference metric does not apply (as in the case of agents sequenced using FSAs). However, when continuous differences can be computed, the hierarchic social entropy measure is utilized.