Anigrafs™: experiments in collective consciousness

Whitman Richards

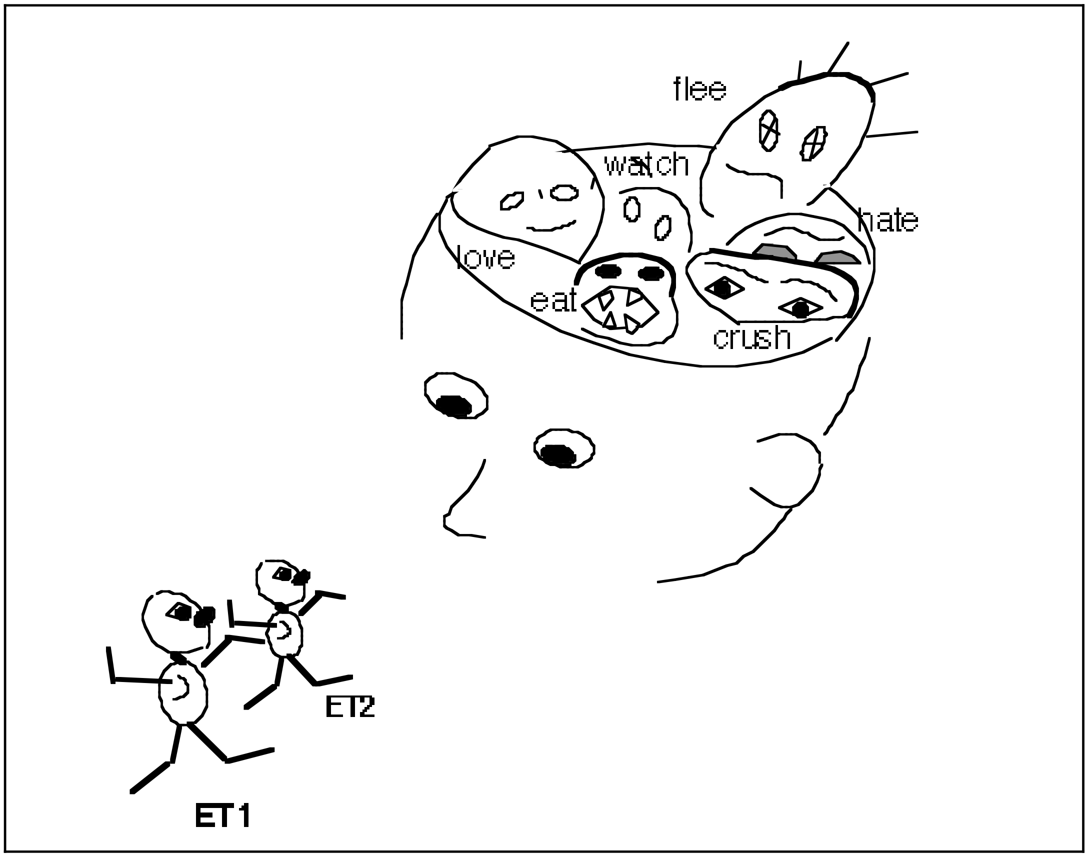

 15 Aug 07

Contents

Acknowledgements: Diana Richards opened new avenues by alerting me to chaotic dynamics in social choice theory. This led to a collaboration that showed how instabilities in choice outcomes could be eliminated if individual preference orderings of alternatives were consistent with a shared global model of the domain. Although in the domain of social choice, these results dovetailed nicely with earlier work with Allan Jepson on perceptual inference. Here also, the key was to assign preference orderings, in this case to possible interpretations of sensory evidence. These posets were rankings of qualitative priors on possible world states. In both domains, the models captured knowledge about relationships and regularities in either the natural or social world. Such regularities reflect tight correlations among alternative observations and events, according to the "Natural Modes" hypothesis suggested many years earlier with Aaron Bobick. The modal correlations appear here in the form of graphical representations, and several new formal results were made possible with the help of Brendan McKay. The integration of these three sets of ideas led to Anigrafs: a framework for information aggregation and choice at many levels, ranging from simple brains to complex societies, such as nation states.

P. Gunkel kindly granted permission to reproduce some charts, and his comments from the perspective of Ideonomy were most insightful. T.J.Purtell designed the Quilt-like phase plot representation for top-cycles. Other collaborators on technical issues were Brendan MacKay, Nick Wormald, Galen Pickard, Sebastian Seung, and Rajesh Kasturirangan. …Many thanks also to the students of Anigrafs who helped debug various proposals and designs. Some of their projects appear in the on-line sub-appendix.

To Waltraud....

......and mentors past and present,

with special affection for Clarence Zener

Preface

About twenty years ago, Valentino Braitenberg wrote a fascinating monograph that described the behavior of some very simple neuronal machines called "Vehicles." The behaviors of these mindless entities seemed so natural, that each type of vehicle was characterized by a name, such as "love" or "hate", "fear" or "aggression". The impulse to assign intentional acts to a machine illustrates that we naturally assume behaviors are typically purposeful and rule-based. This makes sense, for recovering the internal models that govern behaviors allows us to predict actions. The construction of such predictive models is a mental activity that underlies intelligence. Anigrafs is the study of this cognitive world.

To explore this world, we first must recognize that all cognitive models are abstract, like geometries or number theory. Such models belong to a realm quite different from actual physical phenomena. Nevertheless, it is convenient to adopt an approach to the realm of mind that is analogous to the one we use to understand natural phenomena. Hence, just as we consider natural phenomenon to be the consequence of physical entities, we choose to consider mental acts to be the consequence of entities called daemons. These daemons, which are simple mental organisms, will constitute a complex society encapsulated in a machine such as a brain. They will have their own laws and regularities that underlie their collective behavior. Each society of daemons will have its own particular organizational structure and goals for action, depending upon the context. The representational form for this society used here is the *graph*, where the nodes of the graph correspond to the different daemons, and the edges of the graph depict relations between the daemons. To first order, the internal world of the mind is thus represented in the same way that the natural world is represented: as energy bonds or force interactions between component entities, seen at many different space-time scales. The difference, however, is that the force of the interactions in the cognitive world are the result of the communication of information and desires among a society of mental organisms. The structure of these relations among daemons governs the scope and power of the outcomes chosen by their society. The graphical representation for relationships is favored as an artifact that is conceptually very simple, yet has deep and accessible theoretical underpinnings. Because daemons elicit actions considered animate acts, and because they are linked by edges in a graph, the social network is called an Anigraf.

Mental events are usually associated with living creatures. Each creature possesses a complex internal world that has evolved along with the creature's ability for varied kinds of actions. More complex creatures will need more complex models to support these actions. Hence our graphical representation for the mental organisms comprising the internal worlds will proceed from simple to more and more complex graphical forms, following an evolutionary path. As more complex cognitive structures are created, the simpler forms will often become semi-autonomous agents with robotic-like properties, but guided by a hierarchical system of control. Again, this multi-scale structure resembles in many respects how we, as humans, view our own world. This is not a surprise, because all such models are the creations of variety of mental organisms that live within us. These daemons within constitute an internal world that guides our view of the relation between Nature and ourselves.

Whitman Richards
Cambridge, MA

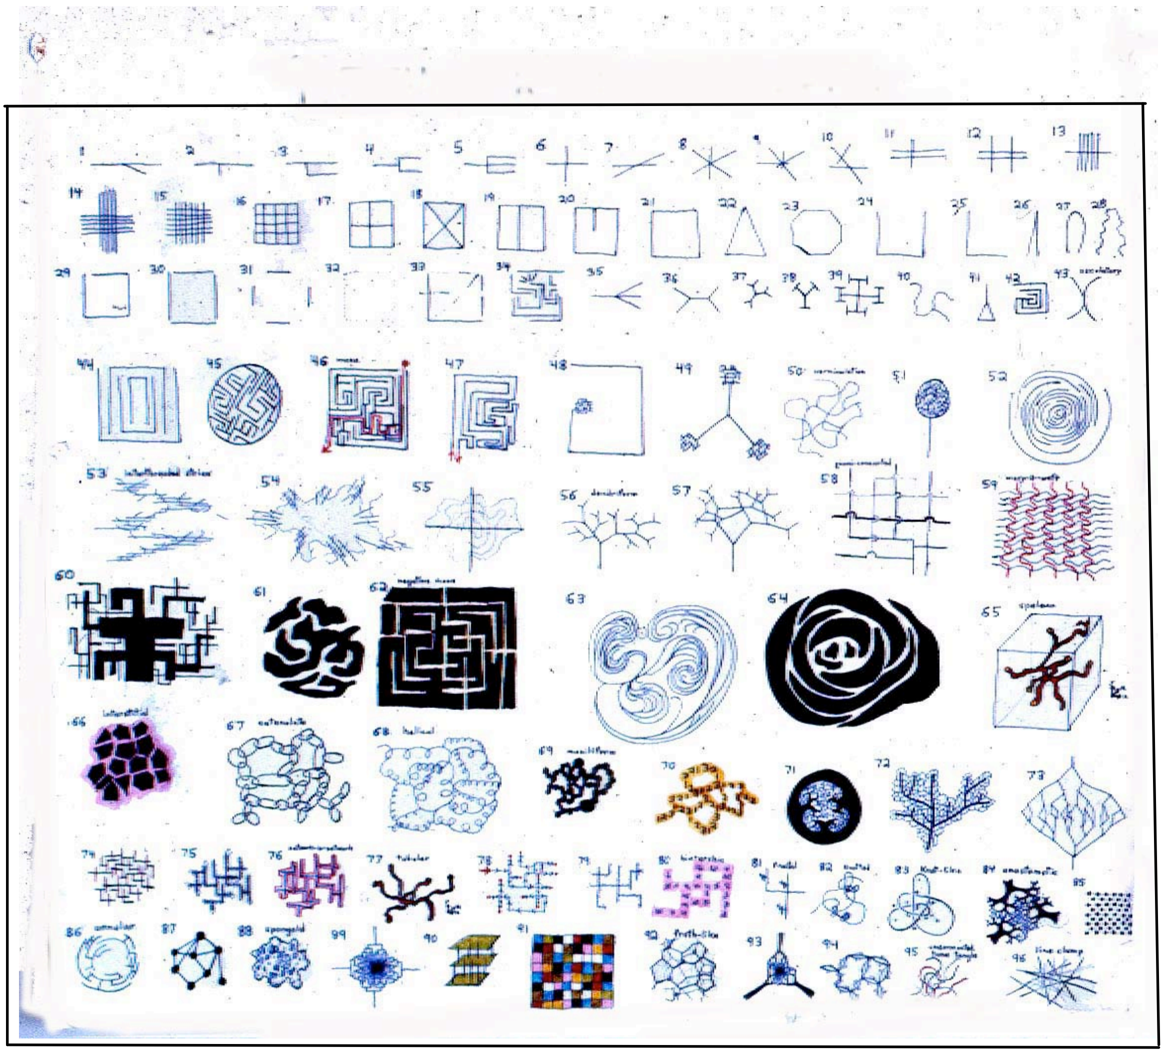

Anigraf Fantasies (from P. Gunkel, the *Ideonomy Project*.)

Part I

Preliminaries
from babble to bargains

Cybernetics offers a view of mind and brain as a collection of mindless agents that act as hierarchical, multi-input controllers. Each agent would be designed to control one variable, such as food cravings, sexual activity, foraging, and for members of bartering societies, costs, benefits, and liabilities. The encapsulated set of such agents would then somehow evaluate trade-offs over many control variables. This is a common view of how brains are organized. However, how should the craving for food be compared with sexual activity, or the need to discuss a problem with a colleague? How should degrees of risk be mapped into a pleasurable experience? Without functions that map one choice into another, the typical cybernetic feedback controller becomes inadequate. Here, we skirt this problem by adopting Kenneth Arrow's argument: choices are described by means of preference orderings without any cardinal significance. This leads us to consider the collection of encapsulated agents or daemons as social-decision makers, guiding complex behaviors by reaching a collective choice based on rankings of preferences. The preliminaries that follow present a framework for aggregating desires, opinions, and needs of a group of mental entities associated with one physical system. Regardless of the complexity of that system, the same procedures for reaching a group consensus will be used in later chapters, because this method can be shown optimal given the constraints of our formulation. Successful biological systems strive for, and often achieve optimality when possible, and we expect the constituent mental organisms to do no less.

24 Jul 07

1.0 Anigraf Abstraction

An Anigraf is a collection of mental organisms called daemons that together form a social system. This system is represented by a graph showing the relations between the desires and actions favored by each daemon. The graphical depiction represents the Anigraf. Figure 1.1 is an illustration. Here we have a puppet with movements of the limbs and body parts controlled by a set of daemons. The daemons pull strings connected to the puppet, translating a mental goal into a physical action. Thus, there is a close coupling between the choices of the group of daemons and the mechanisms for action. It is important, however, not to confuse the Anigraf, which is a mental construct, with the physical acts that the mental agents initiate. The choice of which action will be taken is the result of social agreement, coordinated by communication among the daemons. In the typical world of puppets, such coordination is accomplished by a puppet master. Here, however, an encapsulated society of daemons play this role.

For clarity, the example assumes an anthropomorphic form for the communication network among the daemons, as illustrated by the tree-graph at the top of the figure. However, we could have created a higher-level Anigraf that will decide among possible sequences of actions, such as a set that might be required to pick up the apple. In each case, the graphical form of the Anigraf insures that daemons adjacent in the network will have similar models of the domain, allowing meaningful communication. The Anigraf is thus a social network representing the cognitive world of mental organisms and their relationships to one another. The variety of graphical forms we study range from simple rings of five nodes to very complex scale-free graphs with hundreds of nodes. The thrust of Anigrafs is how these different cognitive designs constrain thought, decision and action, and consequently behavior. More broadly, we wish to understand how behaviors are selected by the mental organisms that occupy creatures as trivial as simple cells to very complex mammals. But first, we need to illustrate how daemons will agree on one choice for action.

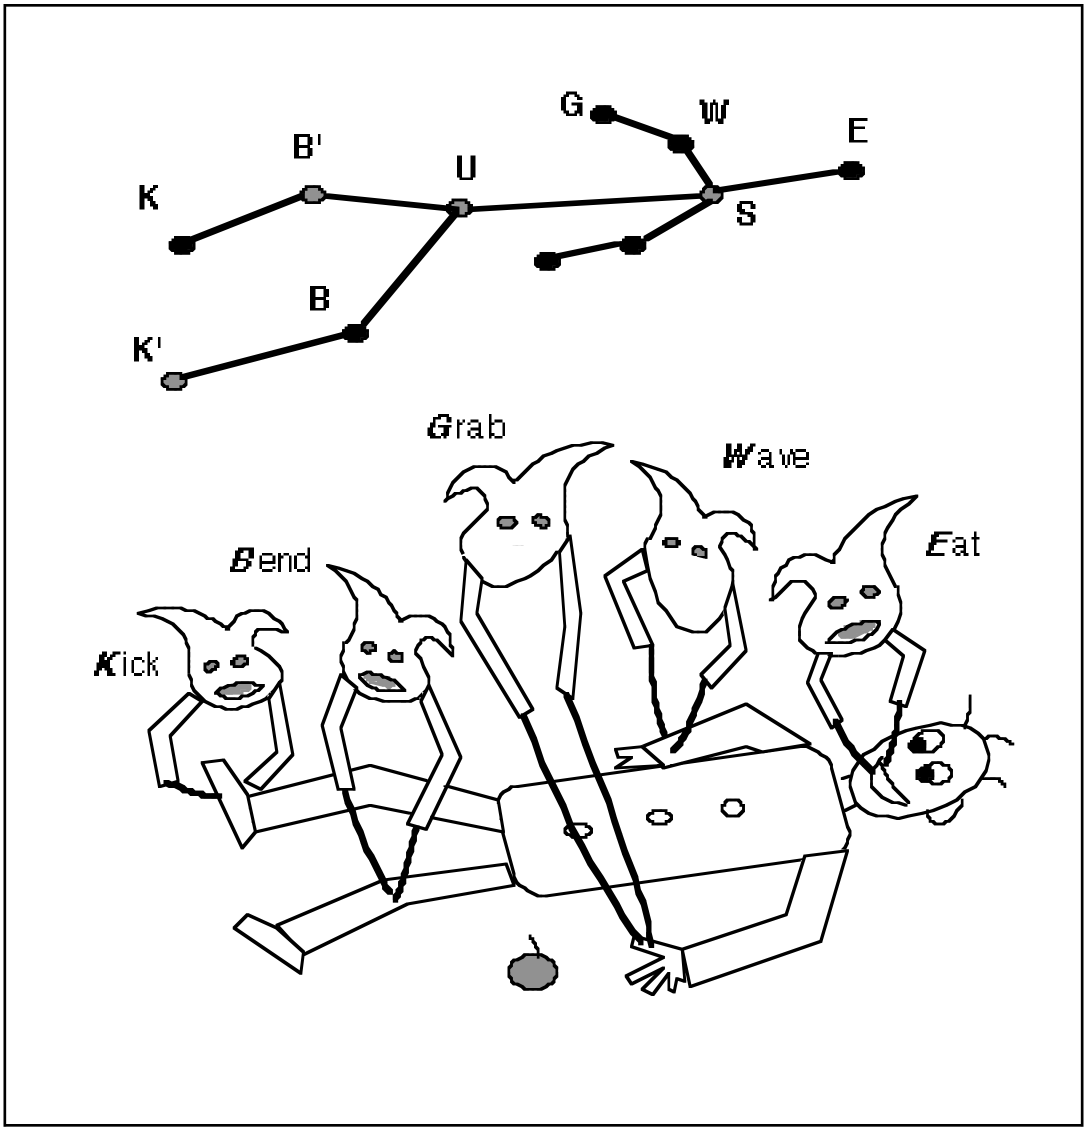

Figure 1.1: an Anigraf (top) showing the abstraction of five mental organisms (daemons) who act as coordinated puppet masters to control actions of a puppet.

2.0 From Pandemonium to Social Contract

Let us activate our puppet. Fig. 1.2 illustrates. In this context, the puppet ponders its environment with uncertainty because of conflicting information held by five daemons. On the one hand, there may be an attractive object ahead seen by sensors associated with a forward looking "***A***pproach" or "***A***ttack"daemon.

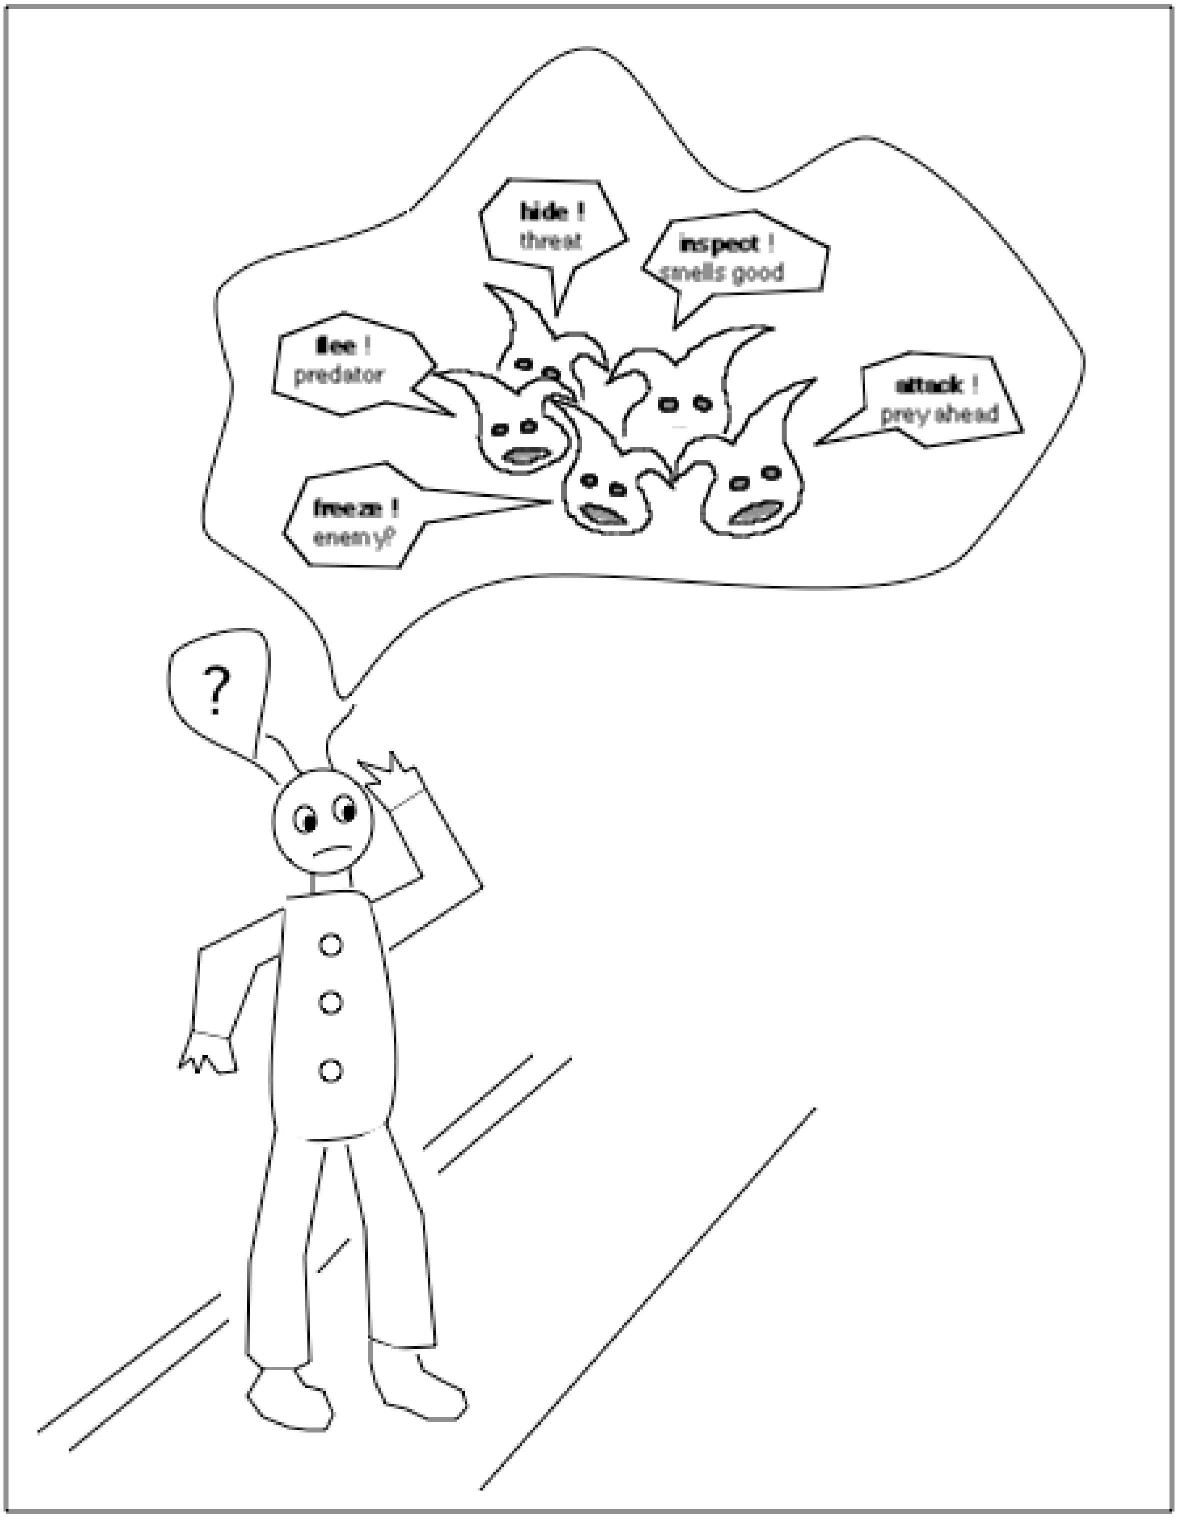

Figure 1.2 A characterization of five daemons actively involved in deciding how the puppet should act in a particular context.

But other daemons, with different vantage points (or sensory organs), may have the sense of a predator lurking elsewhere in the neighborhood. Is the enticing object a lure, put in place by the predator? What should the puppet do? A trivial solution is to let one daemon dictate the option – perhaps a default such as "Free**Z**e." But then much of the information available from other daemons is ignored – perhaps leading to the wrong choice and the demise of our puppet. Anigrafs strive to make use of information held by all mental organisms in their embodied system.

There are many ways to aggregate information, or opinions, or the desires of members of a group. Excluding a dictator, one simple solution is to seek the loudest voice, with this daemon's choice becoming the winner. Let the babble of shouts be given the strengths shown in the left panel of Fig. 1.3. The loudest voice is ***F***lee, with a weight of 10. The collective outcome is then for the vehicle to flee. This is a Plurality or winner-take-all choice aggregation procedure.

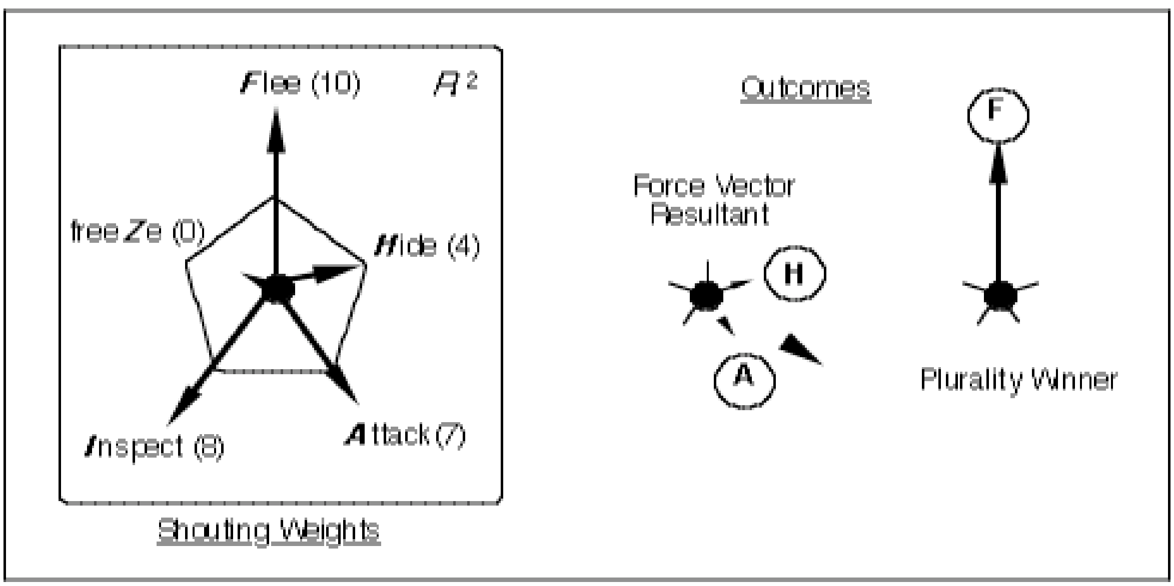

Figure 1.3. Two different procedures for aggregating choices. For illustration, the five daemons have five different vantage points in R^2, as shown on the left. These correspond to directions that the puppet would move to, depending upon which daemon gained control. On the right, the outcomes of two different choice aggregation procedures are shown.

However, another scheme, more typical of physical systems, would be to assign force vectors to each possible action (Fig 1.3, left.) The length of the vectors corresponds to the strength of each daemon's desire. Typically, when such vectors are added, the resultant will point between categories. In our example, the resultant lies between ***A***ttack and ***H***ide (Fig. 1.3, right.) Such half-hearted choices that meld two or more categories are inconsistent with most animal behaviors. Hence, regardless of the complexity of the mental acts entertained, our daemons will make categorical choices. The vector closest to the resultant becomes the sole winner. For our puppet example, the group consensus would now be to ***A***ttack rather than to ***F***lee as the loudest daemon would prefer.

Although very simple and intuitive, the above two methods have serious drawbacks. The weakness of the winner-take-all or Plurality method is its sensitivity to noise – especially if the strongest voice is only a small fraction of the total; and the vector method, aside from its non-categorical nature, assumes that all options are independent, which clearly they are not.

3.0 Intrinsic Knowledge

Returning again to Fig 1.3, note that the five choices were plotted in R^2, as one projection from the five-dimensional space of R^5. This arrangement captures a partial solution to the relations between choices. ***F***lee is the opposite of ***A***ttack, and hence these actions would be negatively correlated in any rational world. Similarly, ***I***nspecting a novel event will require approaching an object, and may be a precursor to an ***A***ttack. ***H***iding has some features in common with standing completely still (Free**Z**e.) Such correlations represent intrinsic knowledge about the how behaviors are seen to be related. They can be used to place choices for action in a two-dimensional metric space (Shepard, 1970), and will provide important constraints for achieving an optimal aggregation of choices, regardless of the complexity of the social system.

To see the impact of shared intrinsic knowledge on group decision-making, let the space of possible actions be related as shown in Fig. 1.4. These relationships are exhibited in two ways, metric (left), and non-metric (right.) The non-metric form links points that fall into adjacent regions. This graphical representation reverals

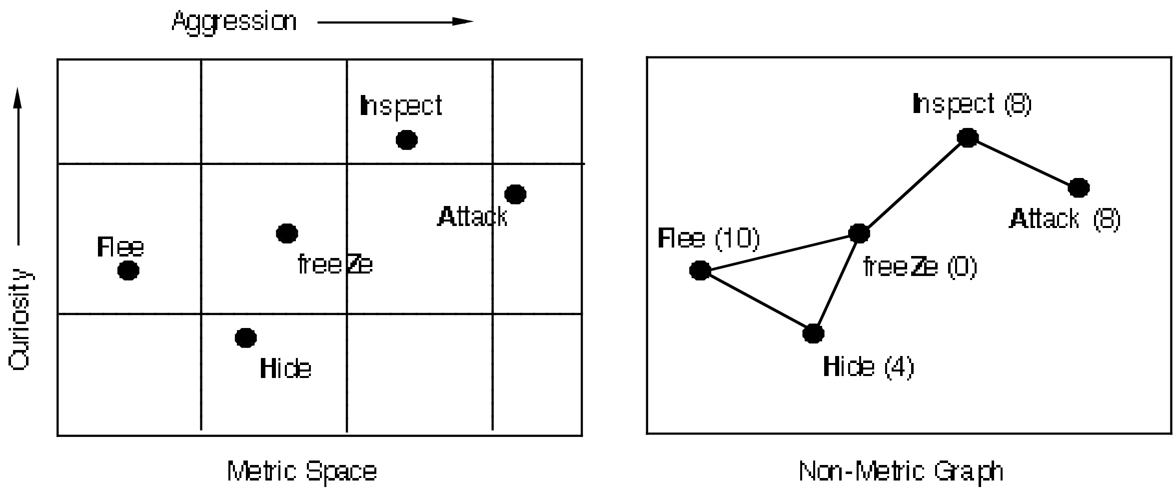

Figure 1.4. Similarity relations between five alternate choices confronting the puppet. Note that the choices are not independent. The graph on the right is a simple example of a shared knowledge structure.

the best second choices for any daemon. If a daemon's first choice is frustrated, then the graph shows the second choices that will at least provide a daemon with some partial satisfaction. The best choice for the society of daemons as a whole will be to maximize the weighted preferences of the group. Hence the graphical representation showing the similarity relations between the choices plays a key role in the daemons' decision-making process. This knowledge is assumed to be intrinsic to the society (in the context), and shared by all members.

4.0 Social Connections: Bartering

The aggregation procedures described earlier for picking winners assume that daemon's choices are independent. There is no need for one daemon to communicate her desires to another. But if some choices are similar or related to others, then this information has value in choosing winners that come closest to maximizing the preferences of the group. How should this information be utilized? Below we begin by identifying types of daemons by the form of their preference orderings on choices. Then we show how such orderings can be used in a tally to find the winning choice for the set of daemons. Finally, we sketch the most optimal aggregation procedure for picking winners (the Condorcet method.)

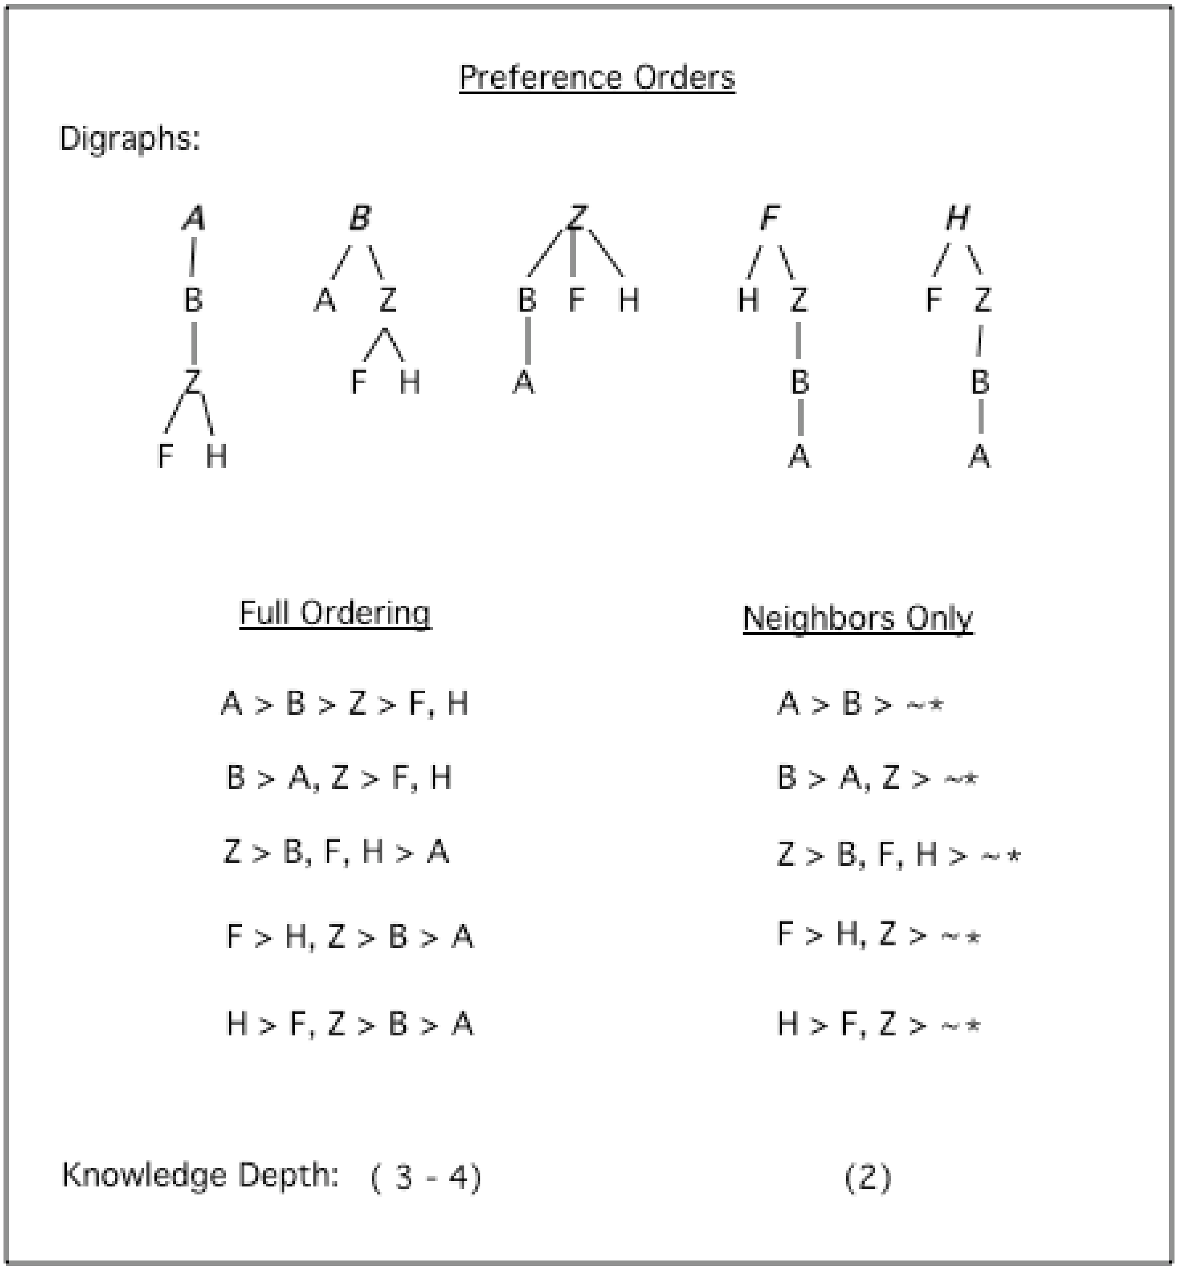

Figure 1.5. Two formats for representing preference orders. Knowledge depth is the maximum number of levels where choices at that level are distinguished by the daemon.

4.1 Preference Orders

Returning to our puppet's five mental organisms, consider again the relations between the choices as shown by the right panel of Fig 1.4. The edges of this graph indicate features shared in common between any two daemons. For example, ***A***ttack and ***I***nspect are forward movements of the puppet, the first being more aggressive than the second. Given this relation, we would expect the ***A***ttack daemon, whose first choice is to attack, would lend support to ***I***nspect over the actions ***H***ide and ***F***lee that take the puppet farther from a perceived threat. The reverse, however, will be true for ***F***lee daemon who favors fleeing and who will most likely regard ***H***iding or free***Z***ing as the best back-up maneuvers. The similarity plots of Fig 1.4 thus make explicit the rank ordering of the alternate choices for any daemon. These are given in Fig 1.5 as a digraph (or more properly a level set) and also in a string form, where > means "preferred to", and ~ indicates "indifference".

Definition: Each daemon is identified by a digraph that places a partial ordering on that daemon's preferences, consistent with the shared intrinsic knowledge about the similarity relations among choices in the context.

4.2 Top Two Choices

We now can engage in modifications of our original social contracts, where daemons can have the hope of at least getting a second choice preference if their first choice is not likely. Let us add to a daemon's shout the voices of all its neighbors on the choice similarity graph (e.g. Fig 1.4.) These neighbors correspond to a daemon's second choice, as shown by the digraphs in Fig. 1.5. (Note that some daemons will have more than one second choice.)

The next step is to include this information in the social choice process. One simple method would be to increase each daemon's level of shouting by adding in the weights of its second choices. For our example, daemon ***A***'s level of shouting will be augmented from 7 to 7+8=15. Similarly, daemon ***Z***'s new shout increases from 0 to 0+8+10+4 = 22. The row labeled "Top Two" in Table 1.1 shows the new aggregated shouting levels for each choice. This procedure picks free**Z**e as the winner. Note that this winner is the first or second choice of all but one of the daemons.

4.3 Borda Count (modified)

A further improvement on the social aggregation method was suggested two centuries ago by Borda (1787.) The advance was to place less weight on second choice, still less on the third, and so forth, averaging over ranks when there were ties. Thus, if we restrict ourselves to considering only first and second choices, a strict Borda Count would give a boost of "2" to the weight of the first choice, a boost of "1" divided among all the second choices and "0" to all the remaining choices. Alternately, in this case, the boosts could be normalized to [1, 0.5, 0] for the first, second and remaining choices. Here, in a modified Borda Method, we apply this normalized boosting to the weights at each level in a daemon's preference ordering. For this modified Borda method, the result is that one-half the combined weight of all the second choices are added to the weight of the daemon's first choice. (See Appendix 3 for details and formal justification.) This weighting is shown by the rank vector in the second column in Table 1.1. The modified Borda winner is ***F***lee with a combined total score of 12.

Table 1.1 First Tally Results

Contract	RankVector	F(10)	H(4)	Z(0)	I(8)	A(7)	Winner
Dictator				***			Z
Vectors	[ ]		*			*	**A**
Plurality	[1, 0, 0]	**10**	4	0	8	7	**F**
Top Two	[1, 1, 0]	14	14	**22**	15	15	**Z**
Borda~	[1, .5, 0]	**12**	9	11	11.5	11	**F**
Condorcet	[ ]						**I**

By now it should be obvious that the procedure used to aggregate the desires of our daemons can have a huge impact on outcomes (Saari, 1998.) Even if information about the choice domain is incorporated into choosing winners, our daemons may still argue over just how second, third, etc. choices should be weighted when votes are tallied. All will agree that second choices should be included in the count, for then there is a clear individual benefit for the majority. But how to settle the rank weighting of second choices? Should the rank vector be [1,0.5,0]; why not [1, 0.25, 0], or even [1,1,0], which is the case for the Top Two tally procedure. Is there a social contract that avoids arguing about rank vectors?

4.4 Condorcet Contract

In 1785 Marquis de Condorcet proposed a scheme that avoided placing weights on lower ranked preferences. His trick was to compare two alternatives at a time – like a tournament – to determine which one is preferred over the other. Now no rank-vector or boosting adjustments on preference levels need be imposed ; each daemon will simply vote for the alternative in the pair that is more desirable. In other words, the daemons picks the member of each pair that is higher ranked in his preference ordering. If one alternative is now found to beat all others in such a pair-wise contest, that alternative is seen as "the fair" social choice for the winner. More importantly, this Condorcet winner can be shown to be the maximum likelihood social choice (Young, 1998; Richards, et al., 2006.) Curiously, as shown in Appendix 1, this method is also an amalgam of the previous Top Two and modified Borda methods.

Table 1.2 sets up a portion of this tournament. The first row gives the voting strengths of the five daemons. The next four rows illustrate how each pair-wise vote is taken. For example, in the first of these rows, F is pitted against I. How will daemon ***F*** vote? Obviously he will choose F over I. Hence alternative F will receive a vote of +10 from daemon ***F***, as shown in the second column of the second row. Next, when daemon ***H*** votes between F and I, because alternative F is nearer in the similarity graph of Fig 1.4, (or equivalently, higher in the digraph of Fig. 1.5), Daemon ***H*** will cast its vote for F, adding another +4. Daemon ***Z***'s position, however, lies equidistant from both F and I, and hence he is indifferent between the two choices, not contributing a vote to either. Daemon ***I***, of course, votes -8 for itself, the negative sign indicating that the vote is cast for the second member of the pair being considered. Finally, for daemon ***A***, choice I is closer to his main

Table 1.2 Pair-wise Condorcet Tally

Pairs	***F***(10)	***H***(4)	***Z***(0)	***I***(8)	***A***(7)	Total	Winner
FvsI	10	4	~	-8	-7	-1	**I**>F
HvsI	10	4	~	-8	-7	-1	**I**>H
ZvsI	10	4	0	-8	-7	-1	**I**>Z
IvsA	10	4	0	8	-7	+15	**I**>A

~ = "indifferent" between the two choices

desire than choice F, hence its vote is cast for I, as shown by the –7 entry. The sum of these entries is –1, indicating that I is the pair-wise winner over F. Following this procedure, the winner for the Condorcet contract is shown to be ***I***nspect, I, which beats all other choices.

The Condorcet pair-wise winner thus has several advantages: first, there is no need for introducing a rank vector when aggregating lower ranked preferences; second, this winner can not be beaten by any counter-proposal, as long as the similarity relationships among the alternatives and weights remain the same; third, use is made of information about choice relations in the domain; and finally, as mentioned, this procedure gives a maximum likelihood choice. Because biological systems strive for optimal choices, the social tallies in subsequent chapters will use this method for aggregating votes.

5.0 Summary

To summarize, an Anigraf is a social network of mental organisms, (or equivalently, daemons) with the following properties:

(1) Daemons' choices are categorical.

(2) There is an intrinsic structure relating choices of the daemons.

(3) This structure is typically depicted as an undirected graph and constitutes the global Anigraf model of the domain.

(4) Each node, or vertex, in the graph corresponds to a daemon with a unique first choice preference for the next state of the social system.

(5) Each edge in the graph indicates that there is a common feature or similarity relation between the preferences held by the linked daemons.

(6) A daemon's ranking of preferences is consistent with the global Anigraf structure and has the form of a digraph.

(7) Each daemon has a say in the next state of the social system, with the strength of his vote a variable. This strength is a weight associated with the particular node in the Anigraf that corresponds to the daemon. Voting power increases (or decreases) with desire, or can be accrued by one daemon representing a group of daemons with identical preference rankings.

(6) Aggregation of daemon's desires will using the pair-wise, Condorcet procedure in order to achieve consensus.

As mentioned earlier, the particular form of the social structure represented by the Anigraf plays a major role in determining whether the collection of daemons will make common-sense choices that most would consider rational. Indeed, rational topologies for Anigrafs will be the major theme of the monograph. Arbitrary connections among a large number of daemons who "go their own way" will be shown to create chaotic behaviors, whereas certain other topologies are guaranteed to yield comfortable aggregate solutions. The key is that the preference relations and belief structures of the constituent mental organisms, or daemons, must be consistent with the knowledge about the world embodied in the Anigraf design. Hence there is a strong coupling between intrinsic knowledge and behavior. The sharing of the same global model among agents is an important component of Anigrafs, and leads to a social Gestalt, which is critical to achieving a collective consciousness.