The Logic of Non-persons

Rohit Parikh

City University of New York
Brooklyn College and CUNY Graduate Center

ICLA-2019, IIT – New Delhi

March 2, 2019

This talk will concentrate on two themes.

1. Two what extent is the reasoning of animals and children logical? What do they think?
2. To what extent can we regard groups: corporations, or political parties, etc. as individuals to whom we can assign goals and beliefs?

In other words, how far can we extend the notion of an individual?

Daniel Dennett's three levels

Dennett defines three levels of abstraction, attained by adopting one of three entirely different "stances", or intellectual strategies: the physical stance; the design stance; and the intentional stance:

The most concrete is the physical stance, the domain of physics and chemistry, which makes predictions from knowledge of the physical constitution of the system and the physical laws that govern its operation; and thus, given a particular set of physical laws and initial conditions, and a particular configuration, a specific future state is predicted (this could also be called the "structure stance"). At this level, we are concerned with such things as mass, energy, velocity, and chemical composition. When we predict where a ball is going to land based on its current trajectory, we are taking the physical stance.

Somewhat more abstract is the design stance, the domain of biology and engineering, which requires no knowledge of the physical constitution or the physical laws that govern a system's operation. Based on an implicit assumption that there is no malfunction in the system, predictions are made from knowledge of the purpose of the system's design (this could also be called the "teleological stance"). At this level, we are concerned with such things as purpose, function and design. When we predict that a bird will fly when it flaps its wings on the basis that wings are made for flying, we are taking the design stance. Likewise, we can understand the bimetallic strip as a particular type of thermometer, not concerning ourselves with the details of how this type of thermometer happens to work.

Most abstract is the intentional stance, the domain of software and minds, which requires no knowledge of either structure or design, and "[clarifies] the logic of mentalistic explanations of behaviour, their predictive power, and their relation to other forms of explanation". Predictions are made on the basis of explanations expressed in terms of meaningful mental states; and, given the task of predicting or explaining the behaviour of a specific agent (a person, animal, corporation, artifact, nation, etc.), it is implicitly assumed that the agent will always act on the basis of its beliefs and desires in order to get precisely what it wants (this could also be called the "folk psychology stance").

At this level, we are concerned with such things as belief, thinking and intent. When we predict that the bird will fly away because it knows the cat is coming and is afraid of getting eaten, we are taking the intentional stance. Another example would be when we predict that Mary will leave the theater and drive to the restaurant because she sees that the movie is over and is hungry.

Note that Dennett left out a crucial condition, *that Mary has a car*. Computer scientists are going to notice this lack because an analysis of algorithms stance implies being aware of the condition that Mary has a car is crucial. Lacking a car she might take a bus, or perhaps walk to a restaurant close to the movie theatre. What she wants and what she believes are not enough. We also need to refer to her *capabilities*.

See Sen, Amartya. "Capabilities, lists, and public reason: continuing the conversation." *Feminist economics* 10.3 (2004): 77-80.

Umwelts

But long before Dennett's *The Intentional Stance*, and Nagel's "What is it like to be a bat?" Jakob von Uexküll carried out a detailed investigation (in the early 20th century) of how animals, children, and we adult humans see the world. The way we *see* the world as contrasted with how the world *is*, is called the umwelt by Uexküll . It is a notion heavily influenced by Immanuel Kant.

"This little monograph does not claim to point the way to a new science. Perhaps it should be called strolls into unfamiliar worlds, worlds strange to us but known to other creatures manifold and varied as the animals themselves . The best time to set out on such an adventure is on a sunny day. The place, a flower strewn meadow humming with insects, fluttering with butterflies. Here we may glimpse the worlds of the lowly dwellers of the meadow. To do so we must first blow in fancy a soap bubble around each creature to represent its own world filled with the perceptions of which it alone knows."

"When we ourselves step into one of these bubbles The familiar meadow is transformed. Many of its colorful features disappear, others no longer belong together but appear in new relationships. A new world comes into being. Through the bubble we see the world of the burrowing worm, of the butterfly, or of the field mouse; the world as it appears to the animals themselves, not as it appears to us. This we may call the phenomenal world or the self-world of the animal."

Jakob von Uexküll, *Forays into the worlds of animals and children,* 1934

How did the world look to him?
Baby Shiva at the age of four months.

And a year later?

Davidson on animals

Neither an infant one week old nor a snail is a rational creature. if the infant survives long enough he will probably become rational while this is not true of the snail....

The difference consists, it is argued, in the having of propositional attitudes such as belief, desire, intention, and shame. This raises the question of how to tell when a creature has propositional attitudes. Snails, we may agree, do not but how about dogs or chimpanzees?...

It is next contended that language is a necessary concommitant of any of the propositional attitudes. This idea is not new, but there seem to be few arguments in its favor in the literature. One is attempted here.

Rational Animals, , *Dialectica*, 1982

Do animals have beliefs?

Norman Malcolm tells this story which is intended to show that dogs think.

Suppose our dog is chasing the neighbor's cat. The latter runs full tilt towards the oak tree but suddenly swerves at the last moment and disappears up a nearby maple. The dog doesn't see this maneuver and arriving at the oak tree he rears up on his hind feet, paws at the trunk as if trying to scale it and barks excitedly into the branches above. We who observe this whole episode from a window say "he thinks that the cat went up the oak tree"

Davidson, *loc cit.*

But how about the dog's supposed belief that the cat went up that oak tree? That oak tree as it happens is the oldest tree in sight. Does the dog think that the cat went up the oldest tree in sight or that the cat went up the same tree it went up the last time the dog chased it? It is hard to make sense of the questions but then it does not seem possible to distinguish between quite different things the dog might be said to believe?

Davidson, *loc cit*

Davidson's claim is that a dog chasing a cat up a tree could not have the belief that there was a cat in the tree. The dog might just have had the belief that a furry animal, or even a funny object, was in the tree..

But the argument proves too much. For by the same token a child who has not had sex education cannot know that it has a mother. .

Surely we do not want to go there.

Perhaps it makes more sense to say that the dog's understanding of the concept cat is a little looser than ours and corresponds to a larger equivalence class in the dog's partition of the world. It includes cats of course but also furry animals and perhaps even furry things which move in a purposeful way.

We will address this issue a bit later.

In the section of A Treatise of Human Nature entitled, "Of the Reason of Animals," Hume argued by analogy that since animals behave in ways that closely resemble the behaviors of human beings that we know to be caused by associations among ideas, animals also behave as a result of forming similar associations among ideas in their minds. Given Hume's definitions of "thought" and "reason," he took this analogical argument to give "incontestable" proof that animals have thought and reason.

Robert Lurz in Animal Minds, *The Internet Encyclopedia of Philosophy*

Dogs for instance have a much better sense of smell and much better hearing than we do. But they are partially color blind and their vision is poorer. Their umwelts are different from ours and they have beliefs and desires and plans for action within their umwelts.

Similarly, the blind character Wally and the deaf character Dave in the movie *See no Evil hear no Evil* have different umwelts from each other and from us.

Thus the umwelt is the semantics (or semiotics) of the agent. If we see this agent as having beliefs and desires (in the BDI sense) then we need to understand what world these beliefs and desires are about. Logics for action and belief need to use the real semantics of such agents. We will offer a path towards formalizing such logics.

And then we can understand what actions will come about from these beliefs and desires.

Is language really necessary?

Shiva in his father's garden last November

Shiva (then 16 m.o.) and I were together in his father's garden and Shiva wanted to go on a swing. but the steps to the swing are a bit steep and I did not think I could keep him safe. So I refused to take him..

On a previous occasion he had cried when I would not do what he wanted. But this time he did not cry. Instead he said, "Daddy!".

I remembered then that his father Vikram had taken him on the swing on the previous day and I called Vikram. Vikram came and took Shiva on the swing. No tears!

Instead of using tears, Shiva used a bit of logic. But he had (then) hardly any language beyond maybe a dozen words.

The sofa and the air

Suppose that Aruna has a sofa in her living room. If you ask her if she knows that she has a sofa in her living room she will say, "Are you crazy? Of course I know." but if you say to her "How many pounds of air are in your apartment?" She would have no idea.(It could be about 750 pounds in a typical apartment)

The sofa is in her apartment and so is the air so why does she know about the one but not the other? Aren't they both part of her world? But the sofa is part of her umwelt and the weight of the air is not.

Here is another example. A dog sees his master from a distance but does not recognize him. But when the master comes closer the dog is very happy, wags his tail and licks the master's hand. What is the difference? Dogs orient themselves in the world by smell more than by sight and a distant master is not recognized.

Uexküll is interested in such questions not only for Aruna and for the dog but also for various creatures like a tick or a fly.

Why does the fly get caught in the spider's web? Because a thread in the web is too fine for the fly's vision. So it does not know that the web is there. Once caught, it knows quite well because it is no longer using its eyes but its sense of touch.

There are certain things that we are all supposed to know like whether there is a sofa in our living room but we do not usually know about the weight of the air, even though it too is in our living room.

Following Kant, Uexküll distinguishes between the actual world and the phenomenal world which varies from creature to creature. The phenomenal world is the umwelt.

Now for us humans, our individual umwelt is supplemented by the community umwelts which include information from the umwelts of others, and also from science. The sun and the moon *look* to us as if they are at the same distance but science *tells* us that the sun is much further.

And we certainly did not send a man to the moon using just the phenomenal world. But animals and young humans tend to act primarily or entirely in terms of their phenomenal worlds.

Agent? Or machine?

A Quote

The mechanists have pieced together the sensory and motor organs of animals, like so many parts of a machine, ignoring their real functions of perceiving and acting, and have gone on to mechanize man himself. According to the behaviorists, man's own sensations and will are mere appearance, to be considered, if at all, only as disturbing static. But we who still hold that our sense organs serve our perceptions, and our motor organs our actions, see in animals as well not only the mechanical structure, but also the operator, who is built into their organs as we are into our bodies.

(Uexküll 1957)

.

Two computer scientists respond

On this basis we shall say that an entity is intelligent if it has an adequate model of the world (including the intellectual world of mathematics, understanding of its own goals and other mental processes), if it is clever enough to answer a wide variety of questions on the basis of this model, if it can get additional information from the external world when required, and can perform such tasks in the external world as its goals demand and its physical abilities permit.

(McCarthy and Hayes, Some philosophical problems from the point of view of AI, 1969)

Two computer scientists respond

On this basis we shall say that an entity is intelligent if it has an adequate model of the world (including the intellectual world of mathematics, understanding of its own goals and other mental processes), if it is clever enough to answer a wide variety of questions on the basis of this model, if it can get additional information from the external world when required, and can perform such tasks in the external world as its goals demand and its physical abilities permit.

(McCarthy and Hayes, Some philosophical problems from the point of view of AI, 1969)

Note that McCarthy and Hayes explicitly refer to abilities which were left out of Dennett's account.

If the organism carries a 'small-scale model' of external reality and of its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise, utilize the knowledge of past events in dealing with the present and future, and in every way to react in a much fuller, safer, and more competent manner to the emergencies which face it.

(Kenneth Craik, *The Nature of Explanation*, 1943: 61)

Umwelts

We can think of the umwelt as a homomorphic image of the real world. And that means that some information is *missing*. In view of this missing information the best action is not always the same as the *apparent* best action.

Umwelts

We can think of the umwelt as a homomorphic image of the real world. And that means that some information is *missing*. In view of this missing information the best action is not always the same as the *apparent* best action.

Now the expected value of the apparent best action increases when more information is received. But in order to receive more information the animal needs to develop tools for that, and they incur a cost, so unless the cost is less than the gain the improvement will not be sought. The fly could have had better eyesight and be caught less often, but that more sophisticated eye would be expensive to maintain.

A question raised by Alfred Russel Wallace was why primitive men had brains nearly as large as ours when they did not have to do complex things like file tax returns. But Steve Pinker suggests that even hunters in primitive tribes use very complex procedures to hunt animals.

Having a large brain enables one to make thought experiments and discover the best action on the spot. Animals and plants may have to go through thousands of years of evolution to make the corresponding discovery. This human advantage has had an unfortunate consequence. Certain species of animals were wiped out when the clever humans entered their domain.

In this context, reconsider Uexküll's account of the life story of a tick. A tick has three perceptions. And three effectors (or actions). The typical tick climbs on a grass blade or something similar and waits.[1]

When a mammal passes under the grass blade, its skin releases butyric acid which the tick detects and it drops onto the mammal. It knows it is a mammal because of the warmth.

Then it moves around in the mammal's skin until it finds a bald spot. It sucks blood and then drops to the ground where it lays its eggs and dies.

[1] Apparently a tick can wait for several years without starving to death.

So the tick needs three perceptions,

1) the sunlight to know which direction is up and so to rise,
2) the smell of butyric acid which tells it when to drop and
3) the feeling of warmth which enables it to know that precious blood is available.

It also has three actions, 1) rising, 2) dropping, 3) sucking blood and then (again) 2) dropping.

The tick can be easily represented by a transducer finite automaton.

It also uses *default reasoning* because it does not (bother to) distinguish between blood and some other warm liquid supplied to it by an experimenter. Under normal circumstances it *is* blood and the tick does not need expensive equipment to distinguish blood from fake blood.

Uexküll has lots of examples of creatures being fooled in this way when the best action in their umwelt is not the best action in the real world

Default reasoning is a rational strategy when we would incur too high a cost to deviate from it. It's cheaper to assume that what you see is what you expect to see.

This idea is reminscent of Kant - we perceive the world based on what is presented to us.

Also related is the notion of *indriya* (sense) in Jainism. Each indriya (like smell or sight) is a homomorphism from the real world onto the phenomenal world. According to Jain doctrine it is a greater sin to kill a creature with more indriyas.

Uexküll is skeptical of the idea that there is the "real world". We shall not follow him or ask the reader to. Rather our representation will assume that *there is a real world* which is perceived imperfectly by every creature, whether a bat or a dog or a child. Thus each creature sees a homomorphism from real world to its personal world.

A little bit of mathematics

Definition

An umwelt U consists of two parts. A homomorphism H (many one mapping) from the actual world to the perceived world. And a set A of possible actions. Thus $U = (H, A)$. In addition each creature has a utility function u, so that $u(a, w) = x$ is the utility of action a performed when the world is w. We will assume that x is a real number.

(In actuality it could be some level of satisfaction *for us humans, or the expected number of progeny for animals).*

$H(w) = H(w')$ means that w, w' result in the same perceptions for a particular creature.

Given a world w, the best action $b(w)$ for the creature is that a which maximizes the expected utility $u(a, w')$ over the set $\{w' | H(w') = H(w)\}$. (There is an implicit probability distribution here which we will not specify). The expected value $E(U)$ of the umwelt U is the expected value of the random variable b.

Definition

Umwelt $U' = (H', A')$ refines *umwelt* $U = (H, A)$ *if*
a) $H'(w) = H'(w') \rightarrow H(w) = H(w')$ *and*
b) $A \subseteq A'$.

Thus H' has more information and more abiliities than H.

Theorem

If U' *refines* U *then* $E(U) \leq E(U')$.

The more you know and the more you can do, the better off you are (with some caveats).

Here is the intuitive idea. Suppose I am driving to New Jersey and can take either the tunnel or the bridge. Normally the tunnel is better as it is closer. But it might be closed.

The procedure if the tunnel is open then take the tunnel else take the bridge has a higher utility than either just take the tunnel or just take the bridge. But that if then else procedure can only be carried out in the refined umwelt where the question about the tunnel has been answered.

Thus it pays to know more and it also pays to have more options for action.

Uninformed agent

Here A and B are incompatible conditions which might obtain. X and Y are possible actions of the agent.

	A	B
Action X	$-100, 25$	10
Action Y	6	$-50, 15$

This is a decision theoretic matrix. In condition A, the agent does not know whether the payoff will be 25 or -100 if action X is performed. Thus if A is true then the best action is Y and if B is true then the best action is X. The utility of the umwelt is $(6 + 10)/2 = 8$.

Better informed agent

In this more detailed table P is an additional condition about which the agent could find out.

	A and P	A and -P	B and P	B and -P
Action X	-100	25	10	10
Action Y	6	6	15	-50

So the best action is (if A&P then Y or else if A&-P then X or if B&P then Y or else if B&-P then X). The utliity of the umwelt is now $56/4 = 14$. Knowing about P has pushed the utility up by 6 and so one could say that the knowledge of P is worth 6 units.

Learning more

Why then does a creature not have a maximal U where H is the identity function and A is enormous?

Because acquiring more information and acquiring more possible actions has a cost and the benefit may not justify the cost.

And for Darwinian creatures which rely on evolution to 'learn,' the entire species has to have the extra sensory ability so that *one* creature may benefit. The cost summed up over the entire species may not be justified by the benefit to one member or a few members of the species.

If I have an umwelt U and I ask a question Q then the H becomes refined to a finer H'. The utility of the new umwelt will be greater but the question will have a cost. To ask the question requires me to make sure that the cost is less than the utility gain.

If I am at a train station and ask the agent what time my train is leaving, I will benefit from the answer.

If I am at a train station and ask the agent what time my train is leaving, I will benefit from the answer.

But if I ask how many dishes are available in the dining car, the agent's rudeness will be too high a price to pay for any answer.

Similarly for an increase in actions. If I am going mountain climbing then it makes sense for me to undergo training so that I have more actions available while on the mountain. But if I am not going mountain climbing then the effort gains me nothing.

Suppose that two different creatures have two different umwelts. For example a man with eyesight but no legs, riding another man without vision but with legs.[2] Or it could be a dog leading a man who is blind. in that case the combined umwelt would be to the benefit of both. What is essential in that case that the umwelts supplement each other and that their utilities align.

[2] Something like this happens in one of the Sinbad stories.

Symbiosis

Consider two creatures with umwelts U ands U' and a common utility[3] function u
Then the two together have joint umwelt U" whose H" is the least upper bound of H and H' and whose action set A" is $A \cup A'$. I.e. $H"(w) = H"(w')$ iff $H(w) = H(w')$ and $H'(w) = H'(w')$.

[3]The utility need not be common but the two utilities can be compatible. See e.g. John Nash's work on The Bargaining Problem, *Econometrica* 1950. Todd Stambaugh, CUNY, has done an interesting sequel to Nash, see *Economics Letters* 1917.

Then the joint umwelt refines both the individual umwelts and (with a reasonable bargain) yields a higher utility for both creatures. This explains why we have cases of symbiosis among animals and massive cooperation among humans. (There is also the issue of *compatible* utilities. A leopard and a deer do not have compatible utilities unless we think of the leopard as having the *job* of keeping the deer herd under control.),

Here is an example. In the ocean, certain species, like shrimps and gobies, will clean fish. They remove parasites, dead tissue, and mucous.

Another example: The relationship between goby fish and shrimp. The shrimp digs a burrow into the sand and both organisms live there. Because the shrimp is almost blind, the goby fish will touch the shrimp when a predator is near.

Animal logic

A tiger watches a deer going towards a bush from the left. Then the deer is not seen any more. And it has not emerged on the other side. So the tiger knows and believes that the deer is behind the bush..

The tiger is inferring the presence of the deer behind the bush, which it does not see, from the previous appearance of the deer to the left of the bush, and from the non-appaearance of the deer to the right of the bush..

The tiger is inferring some variable free sentences which it does not experience, from other variable free sentences which it has experienced.

A general framework

Suppose we are given a first order theory T with plenty of constants and variable free terms. T defines a relation R between finite sets X of variable free sentences and other sets Y of variable free sentences as follows:
$R(X, Y)$ iff $T \cup X \models \phi$ for all $\phi \in Y$.

Clearly R is monotonic in X, in T, and anti-monotonic in Y.

Suppose the tiger's behavior shows awareness of Y on the basis of X.

Does the tiger then believe T?

Not necessarily. There are many such theories which will work. And the tiger may be using some other means to infer Y.

But it can be harmless if we attribute to the tiger such a theory T as long as we are aware that this is merely a *facon de parler*.

Thus it is fine for us to say, "the tiger acts as if it believes T".

Question: For which relations R does there exist a finite first order theory which 'explains' R?.

One could also ask which R are computable in polytime or even in linear time.

Can logic exist without language?

Two kinds of agents considered by AI are as follows.

1. Stimulus response creatures. These are creatures whose reactions are fixed given what they perceive. In AI they are represented by means of a table. And indeed the head of a Turing machine is just that. It sees something on a square and it acts.
2. Creatures with a 'knowledge base'. These are creatures who have some cache of 'facts' which they revise and which they use to infer other facts.

But do these facts need to be expressed in language?

David Lewis showed convincingly that probabilistic conditionals are not propositions and cannot be. but they are something. Can we represent them using some mechanism other than adding to or subtracting from a knowledge base consisting of propositions? Hanoch Ben-Yami suggests that we can. Gilbert Ryle, preceding him by several decades suggested something similar. Perhaps this is a very fruitful direction to go in in but not in this talk.

Conclusion

We have made a start towards formalizing some ideas implicit in Uexküll, Dennett and Nagel as well as others. Such a preliminary effort must leave many loose ends untied. Here are two examples.

Darwin was puzzled by the long and beautiful tail of the peacock. The tail is expensive and makes the bird easier to catch. So why bother? One explanation is that hens like it and the poor peacock has to fall in line. But the tail does not contribute to the peacock's own utility. Only to the expected number of progeny. So there are two utilities involved here. The peacock's own utility and that of its DNA. (And the utiity of the hens who like to see something pretty.) The two can conflict and then the DNA will probably win. But how do these two utilities bargain?

A second technical point is that the actions in the set A not only have a utility but also change the world in some way. It could well be that a *sequence* of actions a_1 and a_2 is what is actually useful. In that case the only benefit of a_1 is to change the world so that a_2 becomes useful.

But these are also issues for a sequel.

Part two

To What Extent is a Group an Individual?

Decision theory for groups

When an individual is about to make a decision she knows what she knows, she (typically) has an idea of what her options are, and she does not need to communicate with herself. I might write a note to my future self but I do not usually write notes to my current self.

But when a group is considering alternatives of action, it needs a certain amount of internal communication. Some facts may be known to everyone, may be common knowledge, or may be known only to a select few in a subgroup. In any case some knowledge structure is crucial.

Similarly, when an action is to be taken, each member of the group needs to do his part. But he also needs to know what his part is and when he should perform it.

This is why an orchestra needs a conductor. Skilled musicians will not perform well without one, although a string quartet does not need one nor does a small group playing Indian classical music.

"The third corollary invokes a concept that linguists, logicians, and economists have called common knowledge, mutual knowledge, and common ground (2, 9, 25–30). In common knowledge, not only does A know x and B know x, but A knows that B knows x, and B knows that A knows x, and A knows that B knows that A knows x, ad infinitum"
[Pinker et al, 2007]

Such considerations may not matter if as with Margaret Gilbert, two people are going for a walk together. If they are together, they can talk to each other or perhaps one can read the other's movements and follow suit. The issue of communication will be simple.

But complex issues will arise with large groups – like armies. And they may arise also with small groups which are separated by distance. If I am to pick you up at the airport so that we can join a family picnic, you and I need to communicate when and where I am to meet you.

Stuart Shapiro[4] considers the issue of coordination in large groups and he points out that the kinds of methods which Gilbert considered will no longer work. There may for instance not be equality among the actors. One of his examples is a group of people travelling by boat to Nova Scotia and to prevent conflicting actions, they elect a captain. She tells them what to do and thus the actions of the group are in harmony.

In this case, the group itself has a goal, namely to get to Nova Scotia in a harmonious and efficient manner. And note that while the captain has the authority for the time being she does need to do a good job or she will not be captian in the next trip. Thus the passengers have authority over her in the long run just as she has over them in the moment. But in the moment they do not have equality, for she has power over them.

[4] Massively Shared Agency, (2011)

.

Even ants need leaders

"Although group living is associated with the range of benefits, the shared occupation of a communal nesting site is also likely to induce cost associated with nest degradation colony growth and parasitism. Consequently group living species form an emigration in which the entire Society relocates to a new nest site. the mechanisms that coordinate nest site selection and colony emigration have been thoroughly studied in ants of the genus Temnothorax and in the honey bee. In colonies of Temnothorax ants, a key stage in colony emigration - disseminating information about the potential nest-sites is organized via stereotyped physical interactions in which a knowledgeable individual leads a naive nestmate to a new nest site in what is called tandem search. " [Richardson et al, 2017]

"In this paper, we present a series of experiments in which colonies of Temnothorax albipennis ants were repeatedly challenged with a difficult consensus-formation taskselecting between two identical nest sites. In order to trace the propagation of information about the available nest sites within the ant population, each ant was tagged with a unique radio-frequency identification (RFID) microchip and the identities of both individuals participating in each tandem run were recorded. By representing each tandem run as a directed link from the leader to the follower, we are able to reconstruct the network of social interactions underpinning the group decision."

"The results are divided into three parts. In the first, we show that tandem running activity is concentrated within a minority with a consistent membership. The second section describes the basic topological features of the tandem recruitment networks formed by the minority that participates in tandem running. Finally, we examine how the role specializations of the leader and follower within each tandem pair, interact to determine the quality of the tandem run."

==========

The influence of the few: a stable oligarchy controls information flow in house-hunting ants Thomas O. Richardson, Charles Mullon, James A. R. Marshall, Nigel R. Franks, Thomas Schlegel http://rspb.royalsocietypublishing.org/content/285/1872/20172726

Also the set of people working together on a project might not be the same as the group proper. Here is an example. Suppose that the US wants to take military action in Syria. In order not to risk too many American lives it decides to use Turkish troops. Turkey agrees for a price, say the removal of sanctions. Then the Turkish troops are not really part of America. They have no interest in seeing to it that America's policy objectives are fulfilled. They merely want their salaries from the Turkish government and obey its directives. Turkey also may not be interested in America's goals. It merely wants the removal of sanctions.

But if Syria wants to deal with a possible US action against it, it will of necessity need to consider the motives and probable actions of Turkey, Russia and Iran. This problem makes the issue of group identity tricky. Are white activists a part of the Republican party or are they independent - and ornery - agents which the Republican party needs to deal with? Are the antifas part of the Democratic party or, again, an independent agent which the party needs to deal with? It is a little like asking whether a guard dog is a member of the family. Perhaps a lapdog is a member of the family but a guard dog is not.

So the "actors" might not coincide with the group and may not share all of its objectives. Each actor has a certain individual goal, his goal, and if that actor's goal is to achieve the goals of the group, then the actor is probably also a member of the group. If the actor is not really interested in the goals of the group, then it is not really a member but merely a hireling who needs to be paid off in some way.

Prashant Parikh in his The Use of Language, CSLI press, looks at the game theoretic issues implicit in the notion of joint action in a group. He appeals to the idea of Nash equilibrium where each agent does the best for himself provided the others are doing their part.

But probably, group membership involves a winning strategy which is in some sense long standing and similar to a repeated game rather than to a one shot game. Robert Aumann relates the amusing story of Reinhard Selten who always carries an umbrella even in a desert like the Negev. Selten does not want to think about whether it is going to rain today. Similarly a person will not ask, "Is it best for me to be a Muslim today?"

Group identity is a long term strategy, perhaps even permanent, and one may wish it to be common knowledge that one is a member of a certain group. (Or one may not if one is a member of the CIA).

But let us put these issues aside for a moment and assume that we have fixed the problem of group membership somehow and are talking about a particular group whose membership we know and understand.

Utilities

When we look at the actions of a group, we might also need to ask, "What does this group want?" The Democratic party wants to win elections. And after the elections are over, it wants to implement as much of its program as is feasible.

It may need to negotiate with another group, the Republican party, whose utilities are mostly opposed but not entirely opposed.

So I will assume that there is such a thing as group utilities, which may change (somewhat) from time to time and that the actions of the group are to meant to further these utilites.

But what makes them group utilities?

One clue is that group utilities can sometimes increase while the individual utilities of some members decrease.

A kamikaze pilot is decreasing his own utility in a dramatic way while increasing the utility of his group. Similarly for a fire fighter.

Game theoretic considerations

If an individual or another group is interacting with a group, then of necessity she or it needs to have some conception of the other as an agent.

If I am negotiating with the IRS about my tax return, I know roughly what they want, what they are likely to do, etc.

It will be quite different if I am negotiating with my employer (say some corporation) about my salary.

In each case I am dealing with a group which I am – at the moment – regarding as an individual.

But the group's utilities may not always be obvious or clear.

Arriving at Group beliefs and objectives

Three people are thinking of taking a train to some destination and it is agreed that they should board the train only if it is safe, quick, and the route is scenic. And here are the votes of the three travelers.

	Safe?	Quick?	Scenic?	Board?
A	No	Yes	Yes	No
B	Yes	No	Yes	No
C	Yes	Yes	No	No
Group	Yes	Yes	Yes	Yes/No

So do they want to board or do they not?

The source for such issues is a paper by Kornhauser and Sager, followed up by Pettit, List, Pauly and several others. But such issues do not always arise. A group does usually manage to take some decision and we note that all three travelers do have the same utility – they differ only in their opinion.

Is language necessary?

In [Tschudin 2001] bottlenose dolphins were trained to associate a tap signal from a trainer on one or two boxes with the location of a fish reward. In the test trials the trainer in the presence of the dolphin baited one of two boxes behind a screen. The screen was removed and the communicator departed. While the communicator was out of sight the boxes were switched in full view of the dolphin. The communicator then returned and tapped a box. On all test trials the dolphins chose the box that was not tapped. What is more, on those trials where the boxes were switched only upon the communicator's return the dolphins consistently chose the box that was tapped. A plausible explanation of the data, Tschudin suggests is that the Dolphins are attributing a mental state of false belief to the communicator and using her (mis)information to solve the task.
[Lurz 2007]

To be sure, many groups, like the Democratic party, do talk. But it could be that the best theory of their behavior is not through what they say but by understanding the logic that they are using to compute their best action.

An utterance is a move in a game and the agent will tend to make the move which is thought to be the best.

See Parikh, Rohit, and Cagil Tademir. "The strategy of campaigning." *Fields of Logic and Computation II*. Springer, Cham, 2015. 253-260.

Groupiness

So it seems that we need some measure of the groupiness of a group.

1. How tight is the communication system?
2. How similar are the utilities of the various members (in the area of concern)?
3. How continuous is the membership over time?

Conclusion

What we need then is a theory by which we can measure the groupiness or tightness of a group. A tighter group G can accomplish more than another group G' even if both have the same capacities.

But this is for the future...

Acknowledgements

We thank Priya Chakraborty, Juliet Floyd, John Greenwood, Peter Godfrey-Smith, Prashant Parikh, Steve Pinker, Vaughan Pratt, Jesse Prinz and Yunqi Xue for useful suggestions.

References

1. Ben-Yami, Hanoch. "Against characterizing mental states as propositional attitudes." *The Philosophical Quarterly*, 47.186 (1997): 84-89.
2. Dennett, Daniel Clement. *The intentional stance*. MIT press, 1989.
3. Jongjin Kim, *Morality as Social Software: Evolutionary, Anti-entropic, and Epistemic Game-theoretic Aspects*, forthcoming
4. John McCarthy and Patrick J. Hayes, Some philosophical problems from the standpoint of artificial intelligence, Stanford Univesity, Computer Science department, 1969
5. Nagel, Thomas. "What is it like to be a bat?." *The philosophical review*, 83.4 (1974): 435-450.
6. Ryle, Gilbert. *The concept of mind*. 1949.
7. Skyrms, Brian. The stag hunt and the evolution of social structure. Cambridge University Press, 2004.
8. Von Uexküll, Jakob. "A stroll through the worlds of animals and men: A picture book of invisible worlds." *Semiotica* 89.4 (1992): 319-391.