2018

Formalizing the Umwelt

Rohit J. Parikh
CUNY Graduate Center

Formalizing the Umwelt

Rohit Parikh

City University of New York
Brooklyn College and CUNY Graduate Center

Unilog 2018, Vichy

June 26, 2018

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, Donald Davidson, *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'

(Malcolm added we would say the dog was barking up the wrong 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.

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*

Umwelts

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.

Dogs for instance have much a 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. Ditto for the behavior of the blind character Wally and the deaf character Dave in the movie *See no Evil hear no Evil*.

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?

Baby 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 had taken him on the swing on the previous Saturday and I called his father, Vikram. Vikram came and took Shiva on the swing.

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

Preamble

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

Dogs for instance have much a 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. Ditto for the behavior of the blind character Wally and the deaf character Dave in the movie *See no Evil hear no Evil*.

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.

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 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.

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 1969)

From Jongjin Kim 2017

The concept of Popperian Creature of Daniel Dennett (1995, 1996) can shed further light on backward induction. Dennett in his Darwin's Dangerous Idea (1995) and Kinds of Minds (1996) discusses an evolutionary hierarchy of intellectual progress. He calls the hierarchy the Tower of Generate-and-Test, where there are five kinds of creatures.

On the ground floor of the tower, Dennett proposes, the inhabitants are (1) Darwinian creatures, organisms who are blindly generated and field-tested, and then only the best designed inhabitants survive by natural selection. (See the discussion on natural selection in chapter 3.) The generation, test, and survival of a Darwinian creature is possible randomly, therefore, luckily. On the upper level, as a subset of Darwinian creatures, there are (2) Skinnerian creatures, referencing to the behaviorist psychologist B. F. Skinner. Skinnerian creatures blindly try different responses to the environment until one response is selected by reinforcement. And next time, unlike Darwinian, the Skinnerian creature will choose the reinforced response as its first choice.

On the next upper, third, floor, there are (3) Popperian creatures, referencing to the philosopher Sir Karl Popper. A Popperian creature can preselect an action from many options before doing it in the outer environment. A Popperian creature has a filter of inner environment, where its many tryouts (i.e., hypotheses) can be safely tested. On the fourth floor, referencing to the psychologist Richard Gregory, there are (4) Gregorian creatures who import mind tools from the outer cultural environment to construct their better inner environments for better generators and testers. On the fifth floor, finally, Dennett proposes, there are (5) creatures like human beings who can use these mind tools and, most of all, language, in the structure of deliberate, foresightful generate-and-test known as science (1995: 380).

Jongjin Kim, 2017

The tick as machine **or** as a being

(1) The tick typically hangs motionless on bush branches. When a mammal passes by closely its skin glands carry perceptual meaning for the tick: the perceptual signs (Merkzeichen) of butyric acid are transformed into a perceptual cue (Merkmal) which triggers eector signs (Wirkzeichen) which are sent to the legs and make them let go so the tick drops onto the mammal, which in turn triggers the eector cue (Wirkmal) of shock.
(2) The tactile cue of hitting the mammal's hair makes the tick move around (to nd its host's skin).
(3) The sensation of the skin's heat triggers the tick's boring response (to drink its host's blood).

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. (Craik 1943: 61)

Umwelts

We can think of the umwelt as a homomorphic image of the real world.[1] 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.

[1] Vaughan Pratt suggested the term *projection.*

.

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 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.[2] 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.

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

So the tick needs three perceptions, the sunlight to know which direction is up and so to rise, the smell of butyric acid which tells it when to drop and the feeling of warmth which enables it to know that precious blood is available.
It also has three actions, rising, dropping, sucking blood and then (again) 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 the 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. And each creature sees a homomorphism from real world to its personal world.

A little bit of mathematics

Definitions

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).

Given a world w, the best action $b(w)$ for the creature is that a which maximizes the expected utility over the set $\{u(a, w') : H(w') = H(w)\}$. (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* $E(U) \leq E(U')$.

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 *tunnel* or *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	5	$-50, 15$

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.

Better informed agent

Here 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	5	5	15	-50

Suppose that P had probability of 0.5 regardless of whether A or B is true. Then the expected value of the best action without knowing about P is 10 if B is true and 5 if A is true. Knowing about P raises these values to 15 and 25 respectively. Thus knowing about P makes things better.

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 some 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.

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.[3] 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.

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

So consider two creatures with umwelts U ands U' and a common[4] utility 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 $H'(w) = H'(w')$.

[4]The utility need not be common but the two utilities can be compatible. See e.g. John Nash's work on the Bargaining Problem.

Then the joint umwelt refines both the individual umwelts and 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 beliefs

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.

Who has feelings?

From Knobe and Prinz, 2008.

To decide between these conflicting hypotheses, we ran a follow-up experiment. All subjects were given a description of an agent that is not in any sense made up of smaller agents but which nonetheless has a physical constitution radically different from our own:

Once there was a powerful sorceress. She came upon an ordinary chair and cast a spell on it that endowed it with a mind. The chair was still just made of wood, but because of the magic spell, it could now think complex thoughts and form elaborate plans. It would make detailed requests to the people around it, and if they didnt do everything just as it wanted, it would start complaining. People used to call it the Enchanted Chair.

Note that this passage ascribes to the chair only states that do not require phenomenal consciousness. (Indeed, it only ascribes states that people would be perfectly happy to ascribe to a corporation.)

The key question now is whether people will automatically conclude that the chair is also capable of having states that require phenomenal consciousness. Subjects were therefore asked the question: Can the Enchanted Chair feel happy or sad?

In addition, all subjects were also given a brief description of the Acme Corporation. They were then asked a question designed to see whether they would ascribe phenomenal states to that corporation, namely: Can Acme Corp. feel happy or sad?

Both answers were given on a scale from 1 to 7. Subjects once again refused to ascribe phenomenal states to the corporation (average rating: 1.8), but they were happy to ascribe phenomenal states to the chair (average rating: 5.6). This difference was statistically significant. (Knobe and Prinz 2008)

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? And 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. The two can conflict and then the DNA will probably win.

A second 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 issues for a sequel.

Acknowledgements

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

References

1. Dennett, Daniel Clement. *The intentional stance*. MIT press, 1989.
2. Jongjin Kim, *Morality as Social Software: Evolutionary, Anti-entropic, and Epistemic Game-theoretic Aspects*, forthcoming
3. Knobe, Joshua, and Jesse Prinz. "Intuitions about consciousness: Experimental studies." *Phenomenology and the cognitive sciences*, 7.1 (2008): 67-83.
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. Von Uexkll, Jakob. "A stroll through the worlds of animals and men: A picture book of invisible worlds." *Semiotica* 89.4 (1992): 319-391.
7. Ziemke, Tom, and Noel E. Sharkey. "A stroll through the worlds of robots and animals: Applying Jakob von Uexkull's theory of meaning to adaptive robots and artificial life." SEMIOTICA 134.1/4 (2001): 701-746.