KRF

Knowledge Representation Framework Project
Unit for Scientific Information and Learning
KTH - Royal Institute of Technology, Stockholm, Sweden

Memo

Erik Sandewall

Ronnie at the Zoo: A Scenario for a Cognitive Autonomous Agent

This series of memos contains research reports and tutorial texts from the project on the Knowledge Representation Framework (KRF).
The present memo can persistently be accessed as follows:

Memo persistent URL:	`http://piex.publ.kth.se/reports/krf/008/`
Date of manuscript:	202009-11-24

Related information can be obtained through the following www sites:

KRF website:	`http://piex.publ.kth.se/krf/`
The author:	`http://www.ida.liu.se/~erisa/`

1 The Scenario: Ronnie in the Zoo

The heroine of this scenario is Ronnie, a young girl who is going to the zoo with her family to celebrate her birthday. Besides Ronnie herself, the company consists of Father and Mother, Aunt and Uncle, Cousin and Bigbrother. The scenario runs in a simulator where there is an autonomous agent for Ronnie whereas the behavior of the other persons, and of the animals, is more or less preprogrammed. The challenge is to define the behavior of Ronnie exhibiting characteristic aspects of intelligence.

The layout of the zoo is a Manhattan-style grid with *roads* in north-south direction and *paths* in east-west direction. Roads number from 1 to 6 and paths from 1 and up to some reasonable limit. Points in the intersections and halfway between two intersections are *waypoints,* and both Ronnie and the others are able to move from one waypoint to the next at each timestep. The two kinds of waypoints are called *crossings* and *midpoints,* respectively.

Each block between roads and paths is divided into four *fields* in a two by two fashion. Each field may contain animals, but only one kind of animals in each field. A person can *watch* a field (and the animals in it) if he or she is in the *crossing* between a road and a path that is adjacent to the field. In other words, a person that is located at a crossing can choose between watching any of the four fields that are at that crossing, and a person at a midpoint can not watch any field. However, not all fields contain animals.

Animals are characterized by what is their *species* and sometimes by other attributes. All the animals in a field are alike, so we do not assign individual names or identities to them. Animals can be characterized as being *awesome, creepy* or *cuddly.* Elephants are awesome, spiders are creepy and kittens are cuddly, for example. It is rare for a species to have more than one such characteristics, but crocodiles are both awesome and creepy, and tiger cubs are both creepy and cuddly.

In the south-west corner of the zoo, around roads 1 and 2 and low-number paths, there is a cafeteria, a beer place, shooting galleries, and the like, and also the toilets. The grown-ups like to stay around there and talk. In the north-east corner there is a motorcycle museum which is the favorite of Bigbrother. The rest of the place consists of fields with animals, and empty fields that only contain parkland, and playgrounds for smaller children.

Ronnie and Cousin move around the zoo fairly freely. Sometimes Ronnie runs into *problems* that require immediate action, but otherwise she pursues a number of *interests,* or standing goals, which are:

- watch awesome animals
- watch cuddly animals
- watch creepy animals
- get something to eat
- get something to drink
- get an icecream cone

Her *level of interest* in any of them is a function of how long time it was since she did it last. She needs the help of a grownup or of Bigbrother to

get eat and drink, since they can pay. She also needs the help of someone (even Cousin can help) in order to see awesome animals, since the fence is so high that she needs someone to lift her, unless there happens to be a stool at the fence which is unusual.

Cousin is also interested in awesome and creepy animals, but he could not care less about the cuddly ones. However, his biggest interest is to be with Bigbrother and watch the motorcycles.

There is a fairly long list of problems, in particular:

- Ronnie is acutely hungry
- Ronnie is actuely thirsty
- Ronnie sees an animal that seems to be sick
- Ronnie cuddles with a cat and gets scratched
- Ronnie cuddles with a dog and gets bitten
- Ronnie gets into a fight with Cousin
- Ronnie falls and hurts her knee which is bleeding
- Ronnie needs to go to the toilet
- Ronnie tries to go back to her parents in the cafeteria area, but can't find the way

More problems can be added. Ronnie has ways of dealing with these problems, but sometimes a second problem comes up when she is reacting to the first one.

Time proceeds in discrete steps in this scenario, and at each timestep Ronnie makes an explicit or implicit decision about what to do next. This decision is in principle done as follows.

- First, obtain up-to-date information about the state of the world around her
- Identify if there is a problem at present, and classify the problem as *urgent* or *not-so-urgent*
- For urgent problems, use a fast method for choosing an action on it, and perform that action
- For non-urgent problems, use a method for choosing an action, but think a little before performing it, in order to determine whether it is a good idea to do it. This thinking takes one timestep
- If there is no problem at hand, then make a choice for what to do next according to the list of interests above. Consider the interest that has been waiting the longest, but take feasibility into account as well.

It can also happen that the parents come and look for Ronnie, for example, because they think she ought to eat something, or it is time to go home. Ronnie may or may not wish to comply. In particular, if she just started doing something for an interest that has waited for very long, or she is on her way to doing that, then she may avoid the parent seeing her. This is

facilitated because the parents call her name when they are looking for her, even before they have seen her. Therefore she has a chance of going in the other direction. If she does want to go joint the parent then she can instead go in that direction.

2 Implementing the Simulator

The scenario world contains a fair amount of detail, so it is useful to have the Leonardo facilities for organizing the world model.

The first thing is to organize the physical structure of the zoo. We introduce composite entities of the form `(crossing: 2 5)` for the crossing between road 2 and path 5, and `(midpoint: 2.5 5)` for midpoints in the obvious way. Fields are represented as e.g. `(field: 2 5 NW)` for the field northwest of the crossing between road 2 and path 5.

An entityfile is set up with entities for those fields that contain animals or other attractions, with attributes for those fields showing what they contain. It is not meaningful to make an entityfile for all crossings and midpoints entities; it is better to have an initialization routine that generates all of them given parameters for the number of roads and paths.

Actions are defined for both Ronnie and the other persons in the scenario. A simulation *stepper* is defined which updates the state of the scenario world with schematically defined actions for the other persons, besides Ronnie, and with occasional problems which are generated somewhat randomly. At each timestep, one first runs the stepper and then runs the routines for Ronnie's decision-making and action.

Ronnie's behavior is defined in the way described above. More specifically this is done using the following routine for each timestep:

- Classify the situation at hand. This is done using a decision tree whose outcome is either an expression describing a problem, or a null object indicating no problem
- If a problem as identified, then select an action, using another decision tree
- If the problem is urgent, execute the action, and then this is all for the present timestep
- If the problem is non-urgent, then use a causal net to check if the proposed action may cause a problem. Since this takes one timestep, do not perform the action, but if no negative effect was seen then keep in mind to do that action in the next timestep
- If there was no problem and an action for a non-urgent problem was selected in the previous timestep, then execute that action
- If there is a current plan for pursuing an interest, then perform the next action in that plan
- In the remaining cases, select the highest priority one of the interests and make a plan for pursuing it. Check the plan for negative consequences, using the causal net and/or prediction using the resolution method. If no negative consequences are found, then decide on the

plan, but do not start executing it - that will happen in the next timestep

This routine is reasonable from the point of view of human intelligent behavior and is good for a start. Some improvements would be appropriate, however. In particular, in those cases where the routine says to perform an action that was decided in response to a non-urgent problem in the previous timestep, it would be appropriate to make a quick check whether some of the new arrived information implies that that action would now lead to a problem. If so, another solution for the problem should be sought. The same applies during the execution of a plan for pursuing an interest.

Also, of course, if no solution can be found for a problem then Ronnie should just ignore it.

A more full-fledged agent should also be able to deal with situations where there are several problems at the same time. However, that is more complicated and will not be addressed in this scenario.