Entertainment Robots – Myth or Reality

Alexander K. Seewald
Austrian Research Institute for Artificial Intelligence
Schottengasse 3, A-1010 Wien, Austria
alexsee@ai.univie.ac.at, alex@seewald.at

Abstract

This paper presents an overview of the current field of entertainment robotics based on experiences as spectator during the RoboCup 1999 and building an experimental entertainment robot based on the LEGO platform with digital color camera and various other sensors. RoboCat is a robot cat prototype that shows cat-like behaviour in the real-life environment of typical households. For behavioural modelling, the Hamsterdam architecture was chosen. While showing that Blumberg's Hamsterdam offers a new programming paradigm to design intelligent entertainment robots, this paper also aims to decide whether or not truly intelligent entertainment robots are as of yet a myth.

Introduction

First, let us consider the commercially active field of entertainment robotics and some current developments in this field. Then we shall clarify the notion of an entertainment robot and introduce an architecture by (Blumberg 1997) to design entertainment robots in a biologically plausible way, starting at an ethogram.

Nowadays strange things are afoot in the area of entertainment robotics. We see a well-known Japanese firm sell as of now 15,000 entertainment robots at a rather high price, which nevertheless only reflects the integrated expensive hardware, with a very effective marketing policy creating a tenfold over-demand[1]. We see upcoming approaches using wireless links to utilize the massive computing resources of today's desktop computers to offer speech and face recognition. We see low-cost kits for robot building which enable the technically inclined to create a large set of toy robots for various purposes. Nonetheless, with all these achievements anyone shall be hard put to find a customer that can tell the difference between a very expensive, state-of-the-art entertainment robot and a cheap clone in terms of what it can and cannot do. So, from point of view of the market, anybody can create an acceptable entertainment robot; it is the marketing that makes the difference. Why then is this considered a scientific question? Of course because we would like to have something that customers CAN differentiate from less able products by its capabilities. This may or may not be reasonable to expect in the near future – only time will tell.

Consider state-of-the-art robot technology the author experienced as spectator at the RoboCup 1999 in Stockholm which featured a Sony legged-robot league. Actually, they resemble cats much more than dogs. Take the following scene: the robot struggles to walk towards the ball, eventually reaches it, lifts his paw, kicks the ball successfully into the goal – however, unfortunately it was the wrong goal, even though the goal colors were blue vs. yellow with an orange ball in front of a dark green background, despite carefully controlled lighting conditions and even black hoods for the human players who had to rearrange the robots about twice per minute since they kept bumping into each other and into the boundaries! It was still most impressing that some robots "died" in such a convincing, life-like way, showing erratic and spastic random movement before they froze and crashed completely. In virtual death they were most alive. Still, these robots offer some promise: it is only a question how far this hardware can go with the right software. Unfortunately only Sony and a few select scientific research groups are writing new programs. As of now, there is no official programming kit, just a kit to edit and create new movement sequences. A short side-look at eBay tells us that there are eight of the robots for sale on eBay, five of which were almost unused. The following quote is by an anonymous seller on eBay who clearly was not satisfied with the return on investment of his AIBO.

> Dad thought he had a great idea, but it didn't work. Now, I'm selling my AIBO because my kids want a real dog...

> While the idea of paying $2,500 for an electronic product that doesn't perform any useful function strikes us as strange now, the fact that the Aibo dog doesn't do anything productive is central to the whole idea of entertainment robots. Sony's goal is not to create a robotic slave that will do your chores, but rather an electronic companion that will make you chuckle with an endearing turn of its head or a playful paw swipe at a ball.

The previous quote by (Buskirk 1999) summarizes it quite succintly: as of now we are unable to create a true robot slave and this will be true for the near and medium-term future. In fact everyone would also be hard put to find a bipedal robot that can walk arbitrary stairs without knowing their height, length and position in advance.

So, while overselling the capabilities of entertainment robots and massive marketing efforts distorts the view people have of entertainment robots, thus making it harder to market an entertainment robot with realistically described capabilities, there still remains the question as to how to design entertainment robots flexibly given the technological restrictions – in order to once be able to

[1] At the time of writing Sony has announced that it intends to serve all orders from now on, discontinuing the so-called AIBO-roulette.

transcendend technological limitations. This paper will consider only intelligent entertainment robots, using the following intelligence definition due to (Steels 1994).

> Behaviour is intelligent if it maximises preservation of the system in its environment.

For entertainment robots, the environment in question is usually a household. Intelligent entertainment robots are thus those who are able to keep their users interested for a long period of time and are not thrown or stowed away, never to be used again, like most other toys.

An intelligent entertainment robot by this paper's definition has to be able to learn new things or be updated so as to remain interesting. Any toy with a fixed repertoire, however complex, is ultimately doomed. Two approaches to this can be seen on the market: robot kits (very time-intensive, but allow a wide variety of robots) and self-learning robots such as the AIBO. The former is still hard to build and program. Even a simple task can prove daunting to any robot, however well designed, leaving the customer with a strong feeling of failure.

Self-learning is a promising approach but the maximum plateau that can be achieved this way is rather low. Giving a robot too much freedom to learn increases the probability of radical screw-up[2], which may or may not be considered funny by the user. It may be noted that Sony's AIBO seems to have twelve possible stages of development and only four distinctive adult stages – definitely not an unique personality for every one!

Alternative methods not yet considered are downloadable personalities and personality toolkits to set various parameters influencing the behaviour in continuous ways. For robots that are already wide-spread such as the AIBO it may be a good idea to offer a programming kit to design completely new personalities to volunteers and set up an exchange board on the internet. This way, many more personalities could be created and exchanged that would be possible if Sony kept the programming interface under check. In any case there are already efforts under way to create such a kit by volunteers. Another point is that a robot that is active for a longer time has more chance of doing something to surprise its owners. To simulate at least a 12h waking, 12h sleeping cycle or enabling the robot to reload himself, e.g. by exchanging an empty accumulator pack with a full one instantly, would be a rather simple way to increase its ability to seem alive.

Summarizing, intelligent robots should have different robot personalities (e.g. by tweaking personality parameters or downloading new personalities), changing personalities over time (continuously, not in discrete steps!) and learning new behaviours or tricks. The following will show that the Hamsterdam architecture (Blumberg 1997) can accomodate all these points.

[2] i.e. of learning something unwanted, useless or just incomprehensible to humans; thus leaving the customer thinking "dumb robot!".

Background

Sony's AIBO is a four-legged robot with a head and binocular vision and was programmed to chase and kick an orange ball, to give paw, express joy, boredom and anger by responding to petting and hitting (1-bit feedback). Additionally they participated in the RoboCup 1998 and continue to participate using various software implementations from three selected universities each year. Sony AIBO combines a 180,000 pixel CCD color camera with a single point infrared distance sensor. The color camera is coupled with the CognaChrome system orginally from MIT which thresholds the image to find blobs of color and returns the position of the biggest blobs. The AIBO changes its personality only in discrete steps, triggered by an internal counter that is checked at boot time. It is not clear what the design paradigm of this robot was, but it probably was Brooks' subsumption architecture (Brooks 1986). The robot can be taught by explicitly programming it new movements but can not learn by example.

NEC's R100 features voice keyword and face recognition, speech output and turns it head when you talk to it. It can control your household applicances such as televisions and lights, notifies you of new emails and reads them. It can also send video emails but no text based ones[3]. The R100 uses a wireless link to a desktop computer and uses offline computation. It is less of an entertainment robot than a personal, keyword-driven voice interface to your desktop computer. The idea of offloading computation wirelessly is a bit dated but quite appropriate here. It is not clear what, if any, personality it has – by definition a slave should not have a personality at all.

RoboCat - Hardware

RoboCat is based on the LEGO platform with EyeBot controller board, digital color camera, bumpers and bend sensors. RoboCat is intended to be a robot cat prototype that shows cat-like behaviour in the real-life environment of a typical household. There were two goals: to create an entertainment robot and to do so using a ethologically plausible model, i.e. Hamsterdam which was previously used for various virtual, computer-simulated creatures[4].

As personality metaphor for RoboCat the cat from Whiskas, known from various television advertisements, is taken to be a prototypically playful, very young and inexperienced cat. This personality was taken as basis for the implementation. During implementation and tweaking of personality parameters, various other personalities appeared, e.g. a paranoically fearful cat that continues to move back from an obstacle upon collision far longer than reasonable – sometimes even until hitting the opposite wall.

[3] This would need speech-independent flawless voice recognition of arbitrary phrases.

[4] most notably Silas T. Dog which could even be taught tricks similar to conditioning in real dogs, see (Blumberg 1997).

A design based on LEGO was chosen because there were no other cheap robot hardware platforms yet available that offered sufficient flexibility for designing a cat-like robot. LEGO is cheap, flexible yet stable and allows fast rebuilding and prototyping. Sensors, actors and other devices can be attached by means of elastic and adhesive tape. It would also have been possible to use a Sony AIBO, but this was meant primarily as low-cost approach.

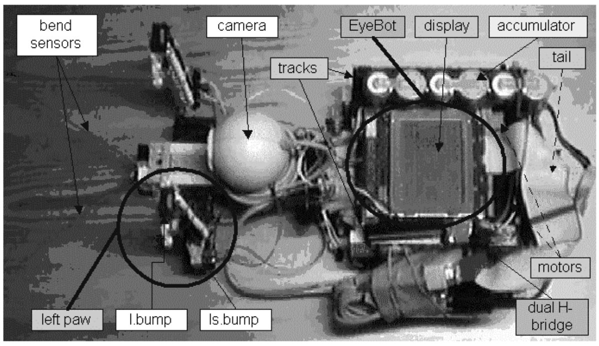

Fig. 1: RoboCat prototype hardware

RoboCat plays with a hard blue ball that is found via image recognition. In every image recognition cycle a 80x50 pixel frame is grabbed and then converted to normalized RGB-scale . Then a box in color space is taken to contain the pixels of the ball, the extent of which has been empirically determined. All pixels of the image are then classified and from those considered to be part of the ball the center of gravity is calculated, the result is the presumed position of the ball.

RoboCat – Software

The behavioural structure and the motivational system was created in four steps.

1. Specification of behaviour
2. Reformulation in robot-centered terms (i.e. as seen from point of view of the robot)
3. Specification of needs
4. Design of the motivational system[5]

First, the desired behaviours will be specified. These will then be reformulated in teleological terms, i.e. specified in a way the robot can understand. Afterwards a system of needs will be designed such that the desired behaviour will emerge. This design will lastly be mapped to the Hamsterdam framework.

Behaviour specification

Simulating cat-like behaviour is a challenging task. The paper (UK Cat Behaviour Working Group 1995) which contains an ethogram of the domestic cat is a good starting point to select interesting behaviours. Let us restrict behaviours to all solitary behaviour patterns and those social behaviour patterns concerning human partners since the author is mainly interested in human-robot observation and interaction. There are still some interesting experiences from cat-robot confrontations although complex social interaction does not emerge.

Since RoboCat has no legs, only one posture element is applicable.

- STAND – Positioned with just four paws in contact with the ground.

Three kinds of TAIL MOVEMENTS were implemented. SLAP was not implemented since the robot's tail is unable to move up or down so it cannot strike the ground.

- SWISH – A cat moves its whole tail rapidly from side to side
- TWITCH – A cat abruptly moves part of its tail from side to side or up and down
- QUIVER – A cat vibrates its tail while raising it vertically

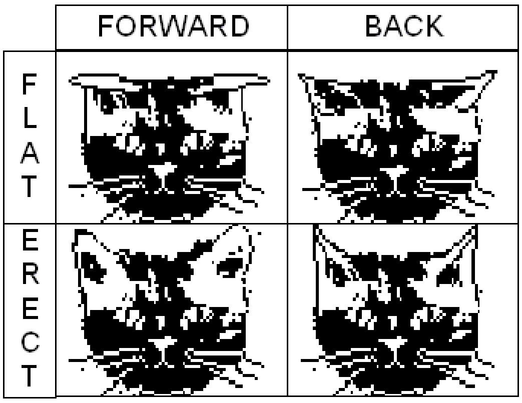

Fig.2: Cat faces with different ear shapes

All four ear positions are implemented as a virtual simplified cat face shown on the integrated LCD screen. The ears can independently be shown BACK vs. FORWARD and FLAT vs. ERECT, yielding four different ear shapes, as shown in Fig. 2. The virtual cat face can also close its eyes and sniff its nose.

- BACK – Ears are held at the rear of the head.
- FORWARD – Ears are held at the front of the head.
- FLAT – A cat flattens its ears to its head such that they tend to lie flush with the top of the head.
- ERECT – A cat points its ears upward.

Four meaow sounds are implemented.

- MIOUW (MEOW) – A cat makes a distinct sound, usually when it is trying to obtain something.
- GROWL – A cat makes a low-pitched rumbling noise.
- YOWL – A cat makes a long drawn-out vocalization.
- PURR – A cat makes a low rhythmical tone from its chest and throat, produced during both exhalation and inhalation.

---

5 My system is based on the Hamsterdam framework from (Blumberg 1997).

The following seven top-level behaviours from Table 1 were modelled. The robot does not build any explicit world model although its internal variables could be read as short-term world model in terms of behavioral tendencies.

Behaviour	Description
WALK	Cat travels fast without obviously investigating its environment.
EXPLORE	Cat travels slowly, sniffing at objects and investigating its surroundings.
SNIFF	Cat raises and twitches its nose, as if to smell.
PLAY(WITH OBJECT)	Cat manipulates an object with its paws in an apparently playful manner.
AVOID	Cat avoids obstacles in its path.
PURR	Cat purrs.
GET ATTENTION	Cat tries to obtain the attention of someone, mainly by MIOUWing.

Table 1 : Behaviour specification

Play and *Avoid* are clearly oriented towards opposite goals. While *Play* will result in interesting interactions, chances are that the robot occasionally tries to manipulate static obstacles, resulting in much pain but not the desired result – the object doesn't move. Therefore the trade-off between the need for pain avoidance and for manipulation has to be resolved by the motivational system.

During testing the original specification the following behaviours also occurred (see Table 2) although they have not been specifically designed. They emerge from interactions between designed behaviours, the robot and the environment. Another type of playing behaviour was also observed, namely moving the ball between paw and whisker.

OBJECT SCRATCH	Cat repeatedly scrapes its extended claws against a rough surface, e.g. wood.
FREEZE	Cat suddenly becomes immobile with body tensed.
RUB OBJECT	Cat rubs its body along the ground or object.

Table 2 : Emergent behaviours

Reformulation in robot-centered terms

Terms like "blue ball", "obstacle", "someone" and "object" that have a more or less definitive meaning for us do not have a definitive meaning for the robot or animals. E.g. what may be an obstacle to an ant, e.g. a pebble, may be no obstacle to a cat and vice versa. Therefore such terms had to be clarified and reformulated in teleological (i.e. robot-centered) terms to make sense for the robot. For example, an obstacle may be described as anything that creates pain (i.e. strong activation of bend sensors or activation of bumper sensors) when the robot moves into it. What the human observer perceives as obstacle may not be perceived as obstacle by the robot and vice versa, although evolution has made a great effort to conceal this disparity between biological entities.

Need	Drive
Pain avoidance	Avoid
Affection	Appetence (human), Get attention, Purr [=*Affection*]
Curiosity Exploration	Wander, Appetence, Sniff [=*Explore*]
Manipulation	Appetence (ball), Manipulation [=*Play*]

Table 3: Which drives fulfill which needs?

Needs

Pain avoidance, *Affection*, *Curiosity*, *Exploration* and *Manipulation* are the main needs of the creature. The author considered teleometaphoresis, i.e. exploiting metaphorical connections between real animals and robots, e.g. hunger as diminishing energy sources, while refraining from simulating immediately useless needs, e.g. the need to drink or to sleep.

Affection may seem a somewhat artificial need. However remember that the robot needs constant supervision and frequent reloading of batteries – so it does need to inspire affection in humans or otherwise it will not survive.

Drives

Drives are defined that satisfy above needs, e.g. a drive to avoid obstacles will satisfy the need to avoid pain, making the robot less likely to experience pain. For an overview of drives and which needs they satisfy, see Table 3.

Design of the motivational system

I used a simplified version of the Hamsterdam architecture as the motivational system. Behaviours are defined that implement the drives. Notice that these are essentially low-level behaviours that may will usually differ from the high-level behaviour definitions we described in the beginning[6]. These behaviours are clustered into a heterarchy similar to Tinbergen's central hierarchy.

Behaviours rely on Internal Variables, which model internal states with autonomous growth/damping and equilibrium points (representing goals and motivations), Level of Interest, which models boredom and relative importance of behaviours, Releasing Mechanisms, which model aspects of the environment which may trigger a certain behaviour and Inhibition as weighted influence from all other behaviours in its behaviour group to determine their own relevance represented as a numeric

6 e.g. the high-level Avoid behaviour was split into four sub-level behaviour for avoiding obstacles encountered from left, right, front and back.

value. All of these influencing factors output numeric values that indicate their strength and are combined to determine the numeric value that indicates the strength of the behaviour.

Internal variables model internal factors that influence the system, such as autonomous growth and damping over time (e.g. hunger), a side-effect of a behaviour being active (ethological term: consummatory behaviour) or a side-effect of a Releasing Mechanism being active (e.g. sensing a snake).

The complete set of internal variable definitions determines the personality of the robot. E.g. if fear is damped less, RoboCat seems more fearful. If aggression is damped less RoboCat is aggressive for a longer timespan, hitting at everything in its path and occasionally at the air. Thus, once a network of interdependent internal variables has been defined, the personality of the robot can be changed by slightly adapting the formulas or even replaced by another set of compatible formulas.

(Blumberg 1997) has demonstrated that by relating changes in internal variables to external events, Hamsterdam is capable of learning new ways to satisfy goals similar to operant conditioning in real dogs. This could principally be implemented for a robot, but was not tested in RoboCat due to severe limits of its hardware architecture. Further details on this system are beyond the scope of this paper, but can be found in (Seewald 1999).

Experiments

Fig. 3 shows a few prototypical cat-cat and human-cat confrontations. The picture on the left shows the author's cat, Bärli, probing its new colleague. The next two pictures shows the two cats playing with a blue ball. On the right, RoboCat shows that it likes its creator, although it probably did not recognize him directly. These pictures were taken from a twelve-minute video[7].

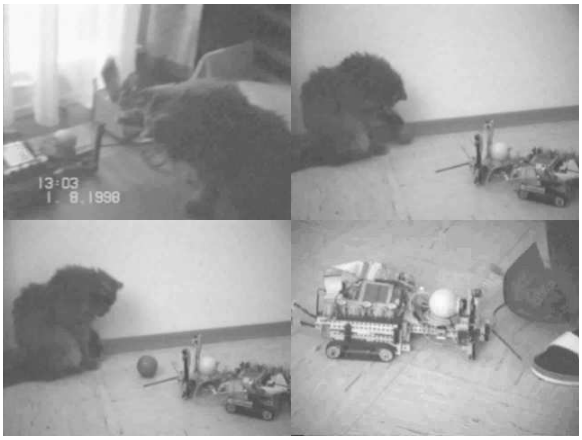

Fig. 3: cat-cat and human-cat confrontations

[7] A low-res 6.4MB AVI is available from www.seewald.at/alex. Also, (Seewald 1999) includes a high-res 340MB AVI on CD-Rom.

Conclusion

I have shown that the Hamsterdam architecture (Blumberg 1997), previously used for computer-simulated creatures, is also applicable to entertainment robotics. Hamsterdam offers a reasonable, practical framework to design behaviour of animal-like entertainment robots based on an ethogram of the species to design appropriate needs, drives and competeing heterachically organized low-level behaviours.

> in the end, i think, this is pandering towards the "AI" perceptions of the masses, who can still be amazed by a 'echo "who are you?"; read idiot; echo "hi, " $idiot;', turning its head to "listen" to people, and so forth. the idea of offloading the computation onto a remote box is brilliant, and should be the way forward, imho, but i think these manufacturers have to get their priorities right.

The final point, emphasized by above anonymous quote on slashdot.org, is that we should aim to make something that is not only an entertainment robot but also a way to explore the limits of robot technology and large-scale robot-robot interaction[8]. The best way to sell entertainment robots is to create an ongoing experience in which we are constantly amazed at what we and other people can teach them or make them do. A start would be to create new motions in a robot not by remotely controlling him, making him into a puppet but just like you teach any other dog – by example. At least until then, intelligent entertainment robots are clearly a myth.

References

Blumberg, B.M. 1997. Old Tricks, New Dogs: Ethology and Interactive Creatures. Ph.D. diss., MIT, MA.

Brooks, R.A. 1986. A Robust Layered Control System for a Mobile Robot. IEEE Journal of Robotics and Automation RA-2, April 1986.

Buskirk, E.V. 1999. CNET.com Review of AIBO. http://www.canada.cnet.com/hardware/0-16332-405-1424637.html?tag=st.co.16332-404-1424637.rev.16332-405-1424637

Seewald, A.K. 1999. A Mobile Robot Toy Cat Controlled by Vision and Motivation. Diploma thesis, Technical University of Vienna, Austria. http://www.seewald.at/alex

Steels, L. 1994. The artificial life roots of artificial intelligence. Artificial Life Journal 1(1):75-110. MIT Press, Cambridge.

UK Cat Behaviour Working Group. 1995. An Ethogram for Behavioural Studies of the Domestic Cat (Felis silvestris catus L.). Universities Federation for Animal Welfare. Potters Bar, England. ISBN 0 900767 901.

[8] How about a get-together of AIBO robots, equipping them with software to talk and exchange data on their interactions with humans?