More Robots in Cages

Exploring interactions between animals and robots.

Marc Böhlen[1]

Abstract

Typically, roboticists have invoked the animal world conceptually as a source of inspiration for finding new approaches to efficient locomotion, perception and intelligent control [Brooks 91], [Hallam, Walker 93], [Aloimonos 97]. The first part of this paper summarizes previous work on designing a robot to share a space with a simple animal. A series of experiments between a mobile robot and three chickens in a cage are described. Techniques are described to mechanically reduce chickens' anxiety towards moving machinery. The second part of this paper places these insights into the wider context of robot design.

Introduction

While the interaction between human beings and machinery, in particular computers (HCI), has received much attention in the past, animal-machine interaction has not. The reasons are obvious. Why would one care? Furthermore, it is difficult to assess how animals perceive things in general, let alone machinery. While the question may be difficult to answer, there could be promising insights. One might be interested in placing a robot in an unusual environment to test its survivability. Additionally, one might want to know how simple animals perceive a non-animated object that can move autonomously. Simple animals understand motion as synonymous with life, and a moving but inanimate object constitutes a novel entity in their world. On a practical level one might like to find design specifications that facilitate the introduction of mobile machinery into industrialized farms, for example.

Design Choices

In order to control the complexities of the problem, the following decisions were made. Chickens were chosen as experimental animals as they are neurologically comparatively simple, cheap and fairly easily maintained in a laboratory setting.[2] A mobile robot was chosen as a representative of machinery in general.

---

[1] School of Art, The Robotics Institute, Carnegie Mellon University, Pittsburgh Pa, bohlen@cs.cmu.edu

[2] The project has been approved by Carnegie Mellon's Research Evaluation Committee with the identification number A3352-01.

The implementation of the design choices proceeded as follows: Three Rhode Island Red chickens and a simple, rugged, can-style mobile robot shared a 6x6ft calcium-sand cage for 60 days. Communication was maintained with the robot through radio modems. A camera mounted directly above the cage maintained feedback. Figure 1 shows a diagram of the experimental setup.

Figure 2 shows a flow diagram of the system components. The first step is the acquisition of an image by the camera. This information is fed into the image analysis module that discerns whether a found blob is a chicken or the robot. The arbitration module decides what the chickens are doing. The architecture module maps, either reactively or by a fuzzy cognitive map, a robot action to the observed chicken action. The communication module, finally, sends the appropriate instructions through the radio-modem to the robot where they are continuously executed.

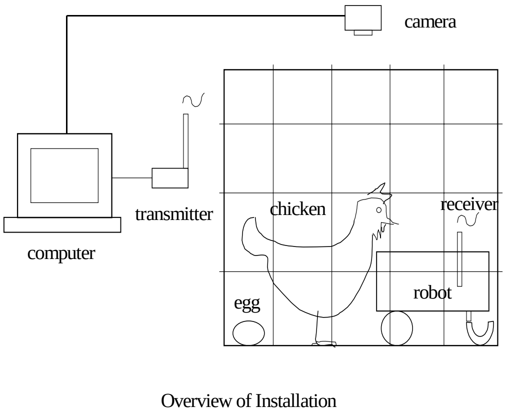

Fig. 1: Overview of Installation

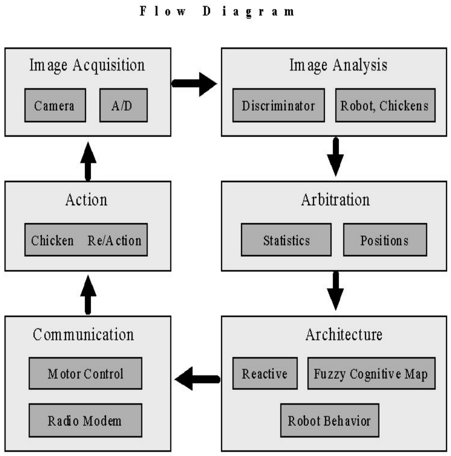

Fig. 2: Flow Diagram

Image Analysis

Each image was then subjected to a blob analysis. In order to discriminate between chicken blobs and the robot blob, filters were applied to each image. The robot filter contained min-max specifications of blob size (area), mean pixel value and aspect ratio. The chicken filter contained the parameters blob size and mean pixel value only. Aspect ratio could not be applied to the chickens as it varied with their activities. If the first filter could not find three chickens, a second filter was called with widened parameter specs. Finding the chickens was a much more error prone task than finding the robot. As the animals were mostly in moderate motion, a ‘lost’ chicken, however, was generally found in the subsequent image.

Navigation

The goal of the navigation module was to (1) be able to keep track of the robot’s position at all times, (2) travel to any specified location in the cage from any part of the cage, (3) travel a specified distance, (4) stop when the robot came within a critical distance of the border of the cage, (5) be able to navigate on a slippery terrain, e.g. full of chicken droppings, (6) avoid bumping into chickens, and (7) recover from a ‘lost’ state, should it occur.

Direction and distance were calculated by maintaining state of the robot over two subsequent images. The slippery terrain made navigation difficult. Feed, water and chicken droppings created an adverse environment for the robot. Random spots of low friction (chicken droppings) in the cage made it impossible to use dead reckoning to navigate the robot. Instead a scheme was implemented by which the position of the robot was compared with the desired position at every instant and corrections were issued according to the progress the robot made in reaching the specified goal. Effectively, this is a servoing technique, using a captured image as feedback, and is generally referred to as *visual servoing*.

Arbitration

Once all the chickens could be reliably found, it was necessary to infer *significance* from the position coordinates of the chickens. While the positions of the chickens and the robot were recorded continuously, only a subset of this information, *significant positions*, was passed into the architecture module.

The first type of significant position was that of a *meaningful position*. For example, it was considered meaningful if the chickens were very close to the robot at any time. The complete set of meaningful positions is: (1) At Robot Home, (2) Close to Robot, (3) By Food Source. To this set of information was added a set of *relational positions.* It included results that indicated whether (1) the chickens were separated out within the cage or (2) all in a group. The last set of positional information was that of *activity*. It discerned whether the chickens were (1) feeding or (2) at rest.

Mapping

The next module in the sequence was designed to map the observations made in the chicken world to that of the robot world.

The mapping concept is based on *mimicking* and *reacting to* the chicken behaviors. If, for example, the chickens were found to be at rest, the robot would rest as well. However, if the chickens were found to be close to the robot, the machine would attempt to get the chickens attention. Figure 3 shows the complete set of correspondences between the chicken world and the robot actions.

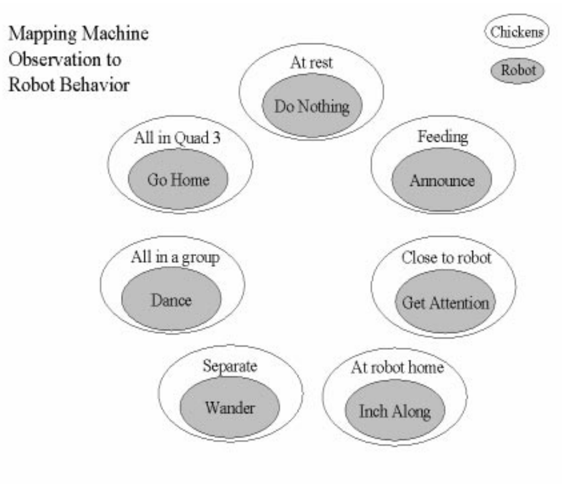

Fig. 3: Machine Observations and Robot Behaviors

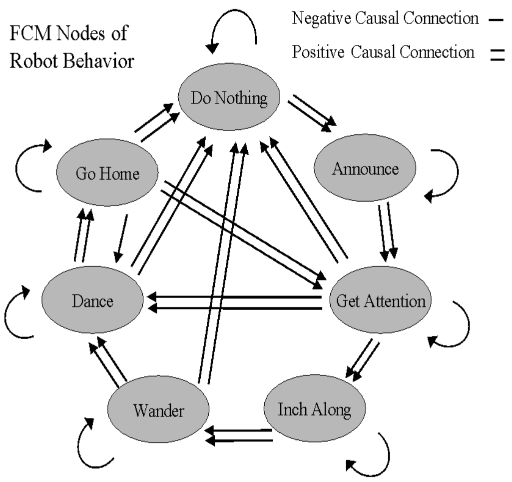

Fig. 4: FCM Nodes of Robot Behavior

Architectures: A Reactive System and a Fuzzy Cognitive Map

The architecture module governs the way the correspondences between the chicken world and the robot were activated. It has two modes of operation. One is a reactive system and the other is a fuzzy cognitive map (FCM).

Figure 4 shows the FCM nodes of the robot actions. A circular transition from low activity (*Do Nothing*) actions to high activity (*Wander*) actions is enforced by strong connections along the nodes. Additionally, intermediately weighted connections at selected points, back to the rest state, allow for a variety of node sequences. A single negative connection between the nodes *Go Home* and *Dance* adds complexity to the higher activity levels.

It is known that FCMs pose a number of problems to the designer. First, one must devise which connections are to be made and second find appropriate weights for them. For a well-defined task such as an environment control scheme, the connections may be very clear. For an interaction scheme such as the one devised between the robot and the chickens it is not. The goal was a design that fostered time-varying, interesting and non-threatening interactions between the robot and the chickens. All connections were chosen to enforce this goal. The setting of the connecting weights was based on trial and error, in accordance with accepted practice for simple FCMs [Kosko 97].

Experiments in Animal-Machine Interaction

The first goal of the animal-robot experiments was to create an environment in which the chickens would share a space with the machine and feel comfortable enough to go about their activities without being threatened.

The first day of interaction with the mobile robot was marked by anxiety and fear on part of the chickens. This is interesting as it confirmed that chickens are confused by the presence of a robotic machine. Apparently they perceive an autonomously moving device as a danger of some sort. It is not clear if they believed the robot to be alive, whether they perceived it as an unknown animal or not. The robot was purposely designed *not* to resemble any kind of animal the chickens might have seen. The reaction was nonetheless clear. The chickens absolutely avoided the machine for an extended period.

In order to attempt to reduce the anxiety of the chickens towards the robot a number of small experiments were performed. One such experiment consisted of moving the robot a few inches, pausing and the then moving again, as opposed to continuous motion. This did not have the desired effect. Driving the robot with reduced speed was helpful in getting closer to the chickens but did nothing to prevent the animals from being scared of the onset of robot movement. It was not quite clear what exactly the chickens were afraid of. In order to test whether the sound of the servo motors was the source of annoyance, the following experiment was performed: The robot was placed in the cage and the motors were powered with short pulses, at first to move the robot only a fraction of an inch, later with no apparent motion at all. The sound of the servos caught the animals' attention immediately. They showed no fear towards the machine as it emitted sound but did not move. This was a surprising result. The next step, then, was to use this fact to reduce the anxiety towards the onset of motion. The servos were pulsed as before, and after a short pause the robot was commanded to move. Interestingly, the birds were not scared of robot motion under this condition. It appeared that the initial sound of the servos caught their attention and that once they directed their attention to the source of the sound, the onset of motion itself was not perceived as threatening. The solution, thus, was to have the robot *announce* itself.

Once it was clear that one could counteract the innate fear of the chickens towards a moving machine, the *announce* solution was made the initial element of various kinds of robot motion. It preceded the *Inch Along* mode, the *Wander* mode and the *Go Home* mode. With this 'trick' it was possible to build a set of robot behaviors to which the chickens showed no anxiety.

Robot Behaviors

The robot behaviors consisted of a set of seven named movement patterns. (1) *Do Nothing* ensured the robot is quiet and at rest. (2) *Announce,* described in detail above, made the robot audibly perceptible to the chickens. (3) *Get Attention* (Fig. 5) subtly alerted the chickens with a flashing set of red and green LEDs and small rhythmic movements. (4) *Inch Along* was a linear motion mimicking the jerky walk of chickens, and (5) *Wander* (Fig. 6) a random walk around the cage. (6) *Dance* was a series of rhythmic and circular movements invoked when the chickens approached the robot and remained very close to it. (7) *Go Home* made the robot return to its defined home base in the lower left corner of the cage.

AMI

Both the reactive architecture and the FCM enabled forms of interactions between the robot and the chickens. For convenience, these shall be termed *Animal-Machine Interactions* (AMI). AMI was enhanced by the intention to make the chickens comfortable in the presence of the robot. This was implemented on two levels. (1) As described above, the behavior primitives were designed to reduce the inherent anxiety the animals had towards a moving machine. (2) Additionally, the robot behaviors were enhanced to include 'politeness' towards the animals. The interaction model was one of *partial hierarchy*. This idea is best explained by examples.

When the robot approached a chicken it stopped and waited for the animal to move out of its path. If the chicken did not move within a few seconds, the robot would attempt to continue moving, but stop again if the chicken had not moved. If the robot moved backwards, it did not care whether there was a chicken behind it. The animals realized this and moved out of the way of the reversing robot more readily than when the robot approached them frontally.

Furthermore, the robot never placed itself for an extended period of time on top of the food source of the chickens. That particular location was 'out of bounds' for the robot. This was an attempt to effectively acknowledge a preferred territory of the chickens. With these enhancements, the previously neutral interactions acquired a flavor of *polite exchange* between the machine and the animals over time. For lack of a better term this enhanced form of interaction shall be termed *Cohabitation*.

Fig. 5: Chicken pecking robot

Fig. 6: Robot inching between two chickens

Experimental Results

The experiments began on February 15 and ended on April 10, 1999. The reactive model was in place the last week of February and the FCM was complete the first week of March. Comparative experiments were performed during four weeks, mornings, afternoons and often evenings. The Reactive model ran during this period for a total of approximately 80 hours and the FCM model for about 60 hours. The following qualitative conclusions are based on recorded video documentation acquired during this time.

(1) It is possible to mechanically reduce the anxiety of a chicken towards a mobile robot. This can be achieved by the following means:

- Announce the intent to engage movement. Do this by either a small meaningless motion followed by a pause or simply an audible sound. Once the chickens direct their attention to the robot, the commencement of continuos motion is perceived as non-threatening.
- Do not move faster than the average speed of the chickens. Speeds approaching their own speed of flight are perceived as highly threatening.

- Pause if the chickens approach the machine or if the machine moves too close to the animals.
- Avoid acceleration. All motion should be continuous.

(2) It is possible to additionally enhance the interactions between the robot and the chickens by implementing a *partial, polite hierarchy* between the participants and choosing an appropriate architecture.

Conclusions

The experiments in AMI described above show the following:

- A mobile robot can be designed to accommodate the innate anxieties chickens have towards moving machinery (*announce-mode).*
- A necessary condition for achieving this goal is to observe and honor the particular habits of the animals, e.g. feeding preferences *(polite hierarchy).*
- Chickens will go about their regular habits with the robot in their vicinity if they perceive the robot as non-threatening *(non-threatening behaviors).*

Future Work

The experiments described above show a machine as a mostly reactive entity amongst simple animals. As opposed to a machine that is tolerated, it would be interesting to attempt to infiltrate the social structure of the chickens and, for example, control the pecking order amongst them. Additionally, architectures including learning and (temporal) planning could be included. Once more results in the interactions with chickens are available, experiments could include neurologically more complex animals. Finally, the experiments show the possibilities and constraints given by a *site specifically situated* robot. The lesson is that the locus and particularities of a site of deployment of a robot are significant both for a good design and for interesting and intelligent interactions with the world.

Site Specific Robotics

Building robots for real worlds is a very different task from designing simulated robots in artificial settings [Brooks 90]. Despite the importance attributed to real world situated robots, much robotics research focuses on robots in engineered, laboratory-style spaces. If robots are to be integrated into everyday life [Agre 97] they must be designed for the most unusual and unreliable spaces. Field Robotics takes on the challenge of unengineered spaces and deploys (mobile) robots into unstructured, natural terrain. The site of deployment, however, is generally considered a nuisance. *Dante II*, for example, a tethered walking robot, able to descend into the interior of a volcano in order to measure ambient atmospheric conditions, was designed to overcome the difficulties of the steep volcanic terrain[3]. The path into the interior of the volcano, however, was considered an obstacle to the 'real' site of interest.

I proposed the term *Site Specific Robotics* [Böhlen 99] to denote an approach of machine design in which design parameters take the particularities of the arena of action into account. The term Site Specific Robotics is borrowed from art theory where *Site Specific Artwork* is understood as an artwork that is particular to a place, situation (and often a time) where it is set. Site specific art is intertwined with the semantics of the place into which it is set and derives its justification from it. An example of site specific artwork is Walter De Maria's "Lightning Field", 1977 [4]. The work (Fig. 7) is to be appreciated only in the site for which it was conceived.

The situated agent concept in robotics should be extended. Not only does it matter that a robot is situated, but also *where and how* a robot is situated. If situatedness and embodiment matter, then a farm is a better test site than a clean room. If the best model of the world is the world itself

---

[3] See http://img.arc.nasa.gov/dante/dante.html

[4] The "Lightning Field" is located in southwestern New Mexico, near Quemado. 400 stainless steel rods are set in a 1mile by 1 kilometer grid in a desolate plain prone to lightning storms.

[Brooks 91], then the complexities of the real world and not a laboratory corridor are the best test bed for a robot intended for intelligent interaction. Furthermore, the goal of designing a general-purpose intelligent robot may not be attainable through the path of general solutions. Horswill [Horswill 93] recast the hard machine vision problem into "light vision", arguing for particular (and fast) machine vision solutions for particular (but complex) situations. The justification for this approach is that, in time, a body of particular solutions can contribute to an understanding of how to design a general solution [Arkin 98]. Social interaction, as often defined in robotics, is but a shadow of what the term promises. It makes sense to study interaction with neurologically simple creatures and learn from these experiences design parameters for (social) interaction [Mataric 94] with more complex beings. In this sense putting more robots into cages and niches maybe a viable iterative method of better understanding what an interactive machine could truly be.

Acknowledgements

This work was supported in part by the Art Department of Carnegie Mellon University, the Robotics Institute and the System's Group. Additional support from Matt Mason, Illah Nourbakhsh and Bryan Rogers made this project possible.

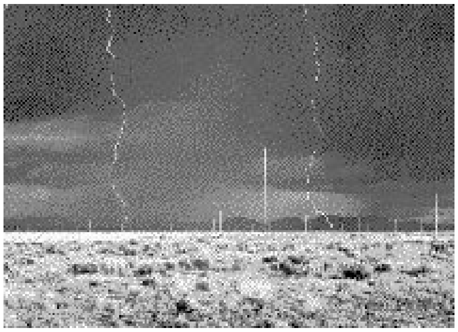

Fig. 7: The Lightning Field

References

[Arkin 98] Arkin, Ronald. "Behavior-Based Robotics". MIT Press, 1998.
[Agre 97] Agre, Philip. "Computation and Human Experience". Cambridge University Press, 1997.
[Aloimonos 97] Aloimonos, Y. "Visual Navigation. From Biological Systems to Unmanned Ground Vehicles". Lawrence Erlbaum Associates, 1997.
[Böhlen 99] Böhlen, Marc. "A Robot in a Cage", IEEE International Symposium on Computational Intelligence in Robotics Automation, Monterey, November 1999.
[Brooks 91] Brooks, Rodney A. "Intelligence Without Reason", MIT AI Lab Memo 1293, 1991.
[Brooks 90] Brooks, Rodney A. "Elephants Don't Play Chess", Robotics and Autonomous Systems Vol. 6, 1990, pp. 3--15.
[Kosko 97] Kosko, B. "Fuzzy Engineering". Simon & Schuster, 1997.
[Hallam, Walker 93] Ashley Walker, John Hallam and David Willshaw, "Bee-havior in a Mobile Robot: The Construction of a Self-Organized Cognitive Map and its use in Robot Navigation within a Complex, Natural Environment", IEEE International Conference on Neural Networks, San Francisco, 1993.
[Horswill 93] Horswill, Ian. "Polly, A Vision-Based Artificial Agent", Proceedings of the AAAI'93, Washington, pp. 824- -829.
[Mataric 94] Mataric, M. "Learning to Behave Socially", From Animals to Animats 3: Proceedings of the Third International Conference of Simulation of Adaptive Behavior, Brighton UK, and pp. 453- -462.
[Rogers 95] Rogers, L. "The Development of Brain and Behavior in the Chicken". Cab International, 1995.