Effective Shared Control in Cooperative Mobility Aids

G. Wasson*, J. Gunderson†, S. Graves* and R. Felder*

{ wasson | gunders | sgraves | rfelder }@virginia.edu
* Medical Automation Research Center, University of Virginia
† Computer Science Department, University of Virginia

Abstract

This paper presents preliminary work on the design of control systems for pedestrian mobility aids for the elderly. The elderly are often restricted in their mobility and must rely on canes, walkers and wheelchairs for locomotion. Restrictions in mobility lead to a loss of independence and autonomy, as well as a decrease in muscular strength.

This paper focuses on design of intelligent wheeled walkers. By allowing the user varying degrees of control, from complete to collaborative, these walkers afford the user with the feeling of control, while helping to increase the ease and safety of their daily travels. The control systems of these walkers differ from those of other mobility aids and mobile robots because they must both assist in mobility and provide balance and support. These functions must be performed in a tight loop with a human whose input may be difficult to predict.

Introduction

The world's elderly population is increasing dramatically. In the US, there are more than 34.8 million seniors over the age of 65. Furthermore, in only 30 years, this number will more than double to 70 million [10]. In Japan, already the nation with the highest percentage of seniors on earth, it is estimated that 1 in 5 people will be seniors within 10 years [17]. At the same time the costs of health care, including caring for the elderly, could rise from its current $1.3 trillion to over $4.0 trillion [6]. If robotic technology can be used to enable the elderly to remain independent, significant costs could be saved and the quality of life would be improved for these people. The attainable cost savings are significant: for every single month that we delay the transition of the elderly population into nursing homes, the US economy saves over $2 billion [1].

Many seniors have mobility impairments that cause a downward trend in their quality of life. Lack of independence and exercise can have dramatic results. Walkers are used more than any mobility aid except the cane [15]. This paper describes our work in designing control systems for wheeled walkers (rollators) to aid in pedestrian mobility of the elderly. We describe various shared control strategies that enable an elderly user and their walker to collaborate on movement, increasing the ease and safety of the user's travels.

Design Issues

The control system of our pedestrian mobility platform is distinct from the control systems typically used in mobile robotics or other assistive devices because it must provide both mobility and balance/support. The difficulty in designing our walker's control system is that the *collaborative* aspect of moving the user must be taken into account.

This collaboration aspect can be viewed as the fusion of two distinct control systems via the physical frame of the walker. One system is the human, providing the motive force and steering the walker towards their goal. The second control system is the steering and braking provided by the walker to avoid obstacles and prevent falls.

These two systems are not independent of each other and care must be taken to insure that errors in the systems do not compound. A walker user anticipates both the movement of the walker frame and the movement of their own body with each step. If the walker's movement differs significantly from expectation, a fall can result. In addition, if the user's expected body motion does not match that produced by their muscles, a fall can occur if the walker is not in a position to allow the user to "catch" their fall. This means that the walker must strike a balance between placing itself in a position anticipated by the user and placing itself in a position that actually supports their musculo-skeletal system.

The control system must take into account more than just the user's balance. Users expect that the walker will maneuver in the way that they push it. The control system must not make the user feel as if the walker is unresponsive or non-obedient. The walker must also help the user reach their destination. Providing the feelings of independence and control to an elderly user, who may have decreased sensor acuity, decreased reaction times, and increased muscle spasticity, is a primary goal of this work. Our success can be measured by the degree to which the control system can help the user do what they "mean to do", as opposed to what their physical input might suggest.

A critical design issue is the ability to integrate the two control systems to provide a smooth sense of shared control. If the autonomous control system in the walker makes abrupt changes to the speed or direction of travel of the walker, it could easily cause the precise problems it is designed to solve. In addition, the operator cannot be expected to enter complex path data into the system while determine their own relevance represented as a numeric 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].

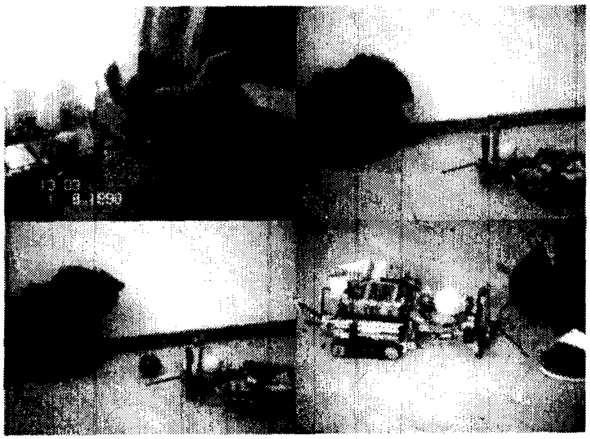

Fig. 3: cat-cat and human-cat confrontations

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.

---

[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.

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

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

---

[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.

much more socially aggressive robot which physically approaches individuals to initiate interaction.

In addition to attracting an audience, a robot must be able to retain one. Museum exhibit designers have tended to make their exhibits more interactive, often even taking on the characteristics of conversation. An exhibit may pose a question requiring the visitor to lift a panel or push a button to hear the correct answer. This is because attention tends to not stay focused through long presentations. By involving the visitors in the exhibit they stay more focused and curious about the information being conveyed.

We have found such techniques for retention to be equally valid for HRI (Nielsen 1993). Chips simply presents long (two minute) video clips at different locations throughout its tour path. As our robots evolved, so did their level of interactivity. Sweetlips includes the human observer in the process of choosing an appropriate tour theme. Joe goes further, answering many different classes of questions and even asking humans limited questions. Adam goes another step, playing trivia games with humans and taking polls. Such an exchange, where both the human and the robot can initiate the next part of the conversation, is essentially *dialogue*.

Because of a robot's particular sensory and effectory strengths, dialogue is multimodal and not necessarily verbal. Thus, while the human may be pushing buttons or using a touch screen, the robot may be responding with spoken words, music, graphics, video, text, physical gestures, and motion.

We learned several lessons from such robotic dialogue design. Firstly, there often will be a crowd of people around the robot, rather than a single person. Together with background noise from the environment, this makes it difficult for some people to hear the robot's responses if they are purely verbal. We therefore ensured that responses are always multimodal, including not only written screen text (e.g. captioning) but also graphics and video content.

Secondly, we found that long presentations, even movies, are guaranteed to drive the audience away. Instead, short responses combined with questions are most effective at extending the conversation. This parallels normal human interaction: the best conversations are dialogues between two people, not lectures. Finally, an aid to increasing the complexity of the dialogue is for the robot to have multiple ways of answering the same question so that it seems less scripted and more spontaneous, and therefore more interesting.

A final lesson learned with respect to HRI involves the psychological effect of creating an anthropomorphic robot. There are strong social rules governing appropriate behavior between humans (though these rules vary between cultures and segments of society), and there are other behavior patterns that people follow when interacting with machines and computers. A robot that looks somewhat human and has a rudimentary personality falls somewhere between these two modes.

The majority of people treat a robot as part human, part machine, clearly following some modified form of human interaction. Often they will treat the robot like a human by default, getting out of its way, and verbally responding to it. If they become interested in some feature of the robot, or want to investigate how it works, however, they will start treating it like a machine, ignoring its requests to move, and standing rudely in its way to see its reaction.

We believe humans use whichever social mode is most convenient for their short term goals. Fortunately, people will also often accommodate a robot that behaves in a fashion that would normally be unacceptable from another human. Since we were not actually in a position to do real social experiments (we had to keep our robot reasonably polite and could not experimentally find the boundary of unacceptable robot behavior) it is difficult to define the extent of this dynamic.

What we were able to experiment with is the robots' displays of emotional state. The main reason for a robot to display emotions is that humans expect and respond to them in somewhat predictable ways. People have a strong anthropomorphic urge and tend to attribute internal state to anything that behaves appropriately. People are also strongly conditioned to react to the emotions displayed by another person. These are powerful tendencies that robots should exploit.

These reactions are entirely behavioral. People cannot discern the true internal state of another human or robot. Their responses are thus entirely dependent upon perceived behavior. Chips and Sweetlips had sophisticated internal mood state machines that would change state over the course of the day, affecting the behavior of the robot. But since the visitors to the museum only interact with the robot for a short period of time, no one noticed these mood changes. Desigining Joe and Adam, we abandoned internal mood representation for a more transparent set of affective reactions to stimuli. On the other hand, if the robots were expected to interact with the same people on a daily basis, the internal moods would once again be useful.

As with the dialogue system, the richer the set of reactions the robot is capable of, the better. For instance, a good interaction model will greet humans in a variety of ways depending on context. If the robot is alone, it should be excited to see someone to interact with. Yet if the robot is busy giving a tour it should politely ask the person to join the tour or, failing that, to please get out of the way so that the tour group can move along.

Even more important than having reactions for all possible interaction contexts, it is critical that the robot's reactions are correct. If the robot begins talking to a wall or to thin air, it looks truly stupid. Just as moving safely through a crowd without hurting anyone is a basic required competence for a mobile robot, so total avoidance of stupid social interactions is a basic competence for a social robot. Generally, no one will notice if the robot fails to react to some indirect stimuli, but they will notice if the robot reacts inappropriately.

In summary, the interactivity of our robots has evolved along four axes: engagement, retention, dialogue, and anthropomorphic/affective qualities. Although this field of research is extremely young, it is already clear that there remains great pliability in the human-robot interaction model: human biases and bigotry regarding robots are not yet strong and fixed. We have an opportunity to design not just robot behavior, but the human behavior that will lead to the most fruitful possible human-robot interaction in the future.

Conclusion

Over the course of the last 2.5 years, we have built four robots, three of which operate on a daily basis with the public, autonomously and without human supervision. While this has been done before (Thrun et al. 2000, Sarcos 2001, Pyxis 2001), our robots are unique in their completely unsupervised free-roaming obstacle avoidance, and in their mission to entertain and inform the general public. We have learned many interesting lessons in attempting to meet the challenges described above; perhaps the most striking is that it actually is possible to deploy robots like these in the public over a long period of time. The robots described above are still running daily, and will hopefully continue to do so for an extended period of time.

In the course of watching the robots change, we have learned many lessons. First of all, it is important to make public robots resilient to physical abuse. People are not afraid to try to damage robots. In fact, they are eager to try to make them malfunction, and especially likely to press large red buttons to see what will happen. Children climb on, kick, and verbally abuse robots. Some fall in love with them. They must be able to handle all of these situations gracefully.

Secondly, when it comes to safety, simplicity in design and paranoia in implementation breeds confidence in deployment. Not surprisingly, once a good and easy to understand system is in place for collision avoidance, it tends not to change.

When robots are placed in public spaces, they must interact with people in such a way that will keep people's attention. The human robot interaction problem is in its infancy. While there have been many experiments in design, few of them have been deployed over the long term, to gauge general public acceptance. Our robots, even though they have been working for quite some time, only scratch the surface of experimentation in this domain. One initial conclusion is that a robot must have an adequate depth of dialog so that a human cannot immediately exhaust the robot's "conversation space," rendering the robot predictable, and therefore uninteresting. But in designing this personality, one must be as conservative as when designing obstacle avoidance code. Making obvious mistakes, such as talking to a potted plant, will cause the robot to be completely dismissed by the audience.

In the domain of autonomy, an approach to design and implementation that implicitly promotes fault-tolerance is important for the long-term survival of a robot. The basic "try again" approach works extremely well since the same code executed twice on a robotic platform will often yield different results. This approach, coupled with the ability of the robots to send pages when they need help, make human supervision refreshingly unnecessary. Even so, there are some types of failures that a robot cannot recover from completely, even if detection of the failure is possible. Drained batteries, burned-out fuses and lightbulbs, and cooked sonar transducers have brought each of the robots described down at various points in time, and the robots simply cannot fix themselves to that degree. Mobile robots still depend on humans for their continuing existence.

Acknowledgements

The original Natural History museum robot concept is due to Jay Apt and Red Whittaker. The following companies provided invaluable help throughout these projects: Nomadic Technologies, Magic Lantern, Maya Design, Redzone Robotics. Thanks also to co-creators of these robots: Judi Bobenage, Ben and Eva Carter, Vinton Coffman, Sebastien Grange, Iacovos Koumas, Ronald Lutz, Roland Meyer, Carolyn O'Brien, Todd Simonds, Alvaro Soto, David White.

References

Nielsen, J. *Usability Engineering*. Boston, MA: Academic Press. 1993.

Nourbakhsh, I., Bobenage, J., Grange, S., Lutz, R., Meyer, R. and Soto, A. An Affective Mobile Educator with a Full-time Job. *Artificial Intelligence*, 114 (1-2), pp. 95-124. October 1999.

Nourbakhsh, I. Property Mapping: A simple technique for mobile robot programming. In *Proceedings of AAAI 2000*. July 2000.

Pyxis Corporation. www.pyxis.com.

Sarcos Corporation. www.sarcos.com

Thrun, S., Beetz, M., Bennewitz, M., Burgard, W., Cremers, A.B., Dellaert, F., Fox, D., Haehnel, D., Rosenberg, C., Roy, N., Schulte, J., and Schulz, D. Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva. *International Journal of Robotics Research,* 19(11):972-999, 2000.