RoboCup for the Mechanically, Athletically and Culturally Challenged

Cindy Marling and David Chelberg
School of Electrical Engineering and Computer Science, Ohio University
Athens, Ohio, 45701, USA, {marling,chelberg}@ohio.edu

Abstract

The Case-Based Reasoning (CBR) research community has recognized the value of learning from failure from the very beginning. Therefore, we can discuss What Went Wrong and Why with impunity. Or without. This paper presents a case-based view of unsuccessful research experience in robotic soccer.

Keywords: failure-driven learning, case-based reasoning, RoboCup

Introduction

> "Cindy Marling is not your typical soccer mom. And Sunday's was not your typical soccer game. As the Ohio University computer science professor watched, one of her players became stuck in a small fence and began spinning its wheels. Literally." (Law 2003)

So said the *Pittsburgh Tribune-Review* on May 5, 2003, the day after the first American Open was held at Carnegie Mellon University. The story continued:

> "By yesterday, just two of Marling's seven robots were functioning. The hardware on the other five gadgets went down. And the shorthanded RoboCats fell to a University of Chile team able to field the full squad of five players."

RoboCup is a Grand Challenge for AI research (Visser & Burkhard 2007). The Ohio University RoboCats, a team in the RoboCup Small-Size League, is a part of this international, interdisciplinary quest (Gillen *et al.* 2002). Maarten Uijt de Haag, from Electrical Engineering, and Jae Lew and Bob Williams, from Mechanical Engineering, joined the AI professors (your authors) in providing faculty leadership for the team. The RoboCats participated in the international tournaments RoboCup 2002, in Fukuoka, Japan, RoboCup 2003, in Padua, Italy, and RoboCup 2004, in Lisbon, Portugal.

We *almost* integrated some Case-Based Reasoning (CBR) into the RoboCats, if you count simulation mode, which unfortunately only counts in the Simulation League (Marling *et al.* 2003). The CBR research paradigm embraces learning from failure, at least for intelligent machines (e.g., see (Hammond 1989)). So, a look at What Went Wrong and Why from a CBR perspective seems apropos. In CBR, knowledge is represented by experiences, or cases, each of which may contain (a) the description of a specific problem; (b) a solution used to solve that problem; and (c) the outcome of applying that solution to the problem (Kolodner 1993). This paper presents several cases from the RoboCup domain in an effort to uncover What Went Wrong and Why.

What Went Wrong

Stuff Blew Up

As duly noted by the *Pittsburgh Tribune-Review*, there were problems with robot *bodies* that frustrated efforts to develop robot *minds*. Early on in the process of designing the robots, we had a number of adverse incidents stemming from a lack of experience in designing and building circuits with high power motors. Our team was and continues to be dominated by undergraduate students who have a lot of enthusiasm, but not a lot of experience, and we get new members on a regular basis who must be trained in the proper way to handle the physical robots.

In one early incident, we were driving a robot around via remote control to get a feel for how fast it could move in various directions, when we heard a loud pop, followed by smoke and a burning smell. We immediately shut down the robot and examined it. One of the chips on our self-designed controller board (called an H-bridge) had completely blown up. Only small pieces were left. We considered the result, and the consensus opinion was that it was just an anomaly... probably, a loose wire had shorted it out. We examined the rest of the robots, making sure there were no other loose wires.

A couple of weeks later, the dean of our college came to see a demo of our robots in action. The demo went very smoothly, with the robots performing well, scoring goals on offense, and defending against human shots on goal. Shortly after the dean left, we heard another, even louder pop, and fire shot out of a robot! The H-bridge on one of the robots used in the demo had blown. After this second incident, we did some more analysis and determined that, much to our surprise, our design was faulty. We changed the H-bridge chips we used and added some protection circuitry to our design. In the several years since this incident, we have had no repeat flames coming out of our "hot" robots.

However, an entirely different system on the robots caused the most extreme explosion, one for which we briefly became famous. Our mechanical engineering students were thinking about how to design the kicker for our robots. Two main ideas were considered. The first design involved using an electrical solenoid (like a pin-ball flipper) to impact the ball, and send it towards the goal. The other design was rather innovative in that it involved using the innards of a paint-ball gun to power the kicker. A paint-ball gun uses compressed gas (CO_2) to expand into a cylinder that impacts the paint ball and sends it out of the gun. Our design used a similar mechanism to push a metal plate towards the ball we wanted to kick. The benefit of the compressed air design is that it uses almost no electrical energy, leaving more power from the batteries available for moving our robots. No other team in the Small-Size League had ever successfully implemented such a system, as it is very hard to fit all the required hoses and CO_2 cartridges into the available space. Several teams in the larger leagues had used rather large compressed air systems to kick full-sized soccer balls at very high velocities.

In any event, we learned one reason why more teams had not pursued our solution to designing kicking systems at RoboCup 2003, in Padua, Italy. Whenever we replaced a spent CO_2 cartridge for our kicker, the procedure called for making sure the system was charged. There is a pin that must puncture the CO_2 cartridge so that the gas can flow out of the metal cylinder it comes in. The students' low-tech solution to determining this was to hold the robot up to their ears. If they heard a hiss, then the system was charged. During a practice in Padua, when a student did this, a tremendous explosion was heard throughout the entire venue. People came running to see what had happened. It turned out that one of the tubes carrying the high-pressure gas towards the release valve had become kinked and ruptured. Luckily, no one was injured, and we changed the procedure to more carefully install new cartridges, inspecting all of the tubes and hoses before each change.

The CO_2 cartridges themselves rendered our kicker design prone to a variety of failures. One problem we never thought of during the design phase was the illegality of the CO_2 cartridges in many countries of the world. In many countries, CO_2 cartridges are not available, or are not available in a form that is useable by our design. Also, airlines are unlikely to look kindly on individuals traveling with hundreds of cylindrical metal cartridges under high pressure. It seems that airport security takes a similarly dim view of these cartridges, which resemble large bullets or small bombs. Due to these factors, when we traveled to Italy, we had to find a solution that didn't involve our taking the cartridges with us. The solution we came up with involved getting a company that had the nearest source of the cartridges (in England) to send a package to the competition site in advance of our arrival. This led to even more complications in terms of customs, international payments, and getting clearance for the package at the venue. Upon our arrival at the venue, we found that no one knew anything about our package, and, of course, everyone only spoke Italian! After getting some help with translation, we finally managed to get them to let us look around a warehouse the size of a large building supply store. Luckily, one of our team members spotted a package labeled with our team name, which just happened to be the correct package. Getting our leftover CO_2 cartridges home again posed another challenge. This time, we sent our most innocuous looking team member to the post office with a cardboard box full of cartridges. We declared on the customs form that the box contained "robot parts," which, of course, it did. The package and its contents arrived safely home in Ohio shortly thereafter. For all of these reasons, our current robots now use electrically powered solenoids for kicking.

The Humans Couldn't Play Soccer, Either

Soccer is not an Ohio tradition, so very few people working on the RoboCats had ever played soccer before in their lives. Clearly, it is easier to imbue robots with intelligence when humans possess intelligence in the domain of interest. To overcome this problem, a series of Human Soccer games was held in the student recreation center for undergraduate, graduate and faculty RoboCats team members. As can be seen in Figure 1, this was highly illuminating. In fact, these games became quite popular, and might have had a successful outcome, had it not been for faulty vision systems. The machine vision system had been left behind in the lab, and human vision was such that the two AI professors (your authors) failed to detect each other as oncoming obstacles. These professors simultaneously kicked the ball full force in opposite directions, leaving the smaller professor with a broken ankle, the larger professor with a fearsome reputation, and twenty students with a deeper understanding of why robots, not humans, should be deployed in hazardous situations.

Figure 1: Professor Demonstrates Ball Control

Cultures Clashed

Cultures collided in three different ways: internationally, across disciplines, and between jocks and nerds. The first problem was that very few RoboCats team members had ever been out of the United States. Student team members did not have passports, speak foreign languages, or voluntarily eat anything unavailable at McDonald's, Pizza Hut, or Kentucky Fried Chicken (KFC). The initial solution was to qualify for international competition, obtain passports, buy plane tickets, and take the students abroad. The expectation was that students would be culturally enriched by being immersed in a culturally rich environment. Expectation failures occurred because: (a) the students were happy to stay indoors working on robots day and night (See Figure 2); and (b) students discovered American fast food franchises wherever they went. (See Figure 3.)

Figure 2: Who Us? Leave the Venue?

Figure 3: The KFC Franchise in Fukuoka, Japan

The repair strategy used in Padua was to declare a ban on American food and to bring fresh Italian breads, cheeses and lunch meats into the venue. After the RoboCats failed to progress past the preliminary rounds (see Figure 4), students were willing to leave the venue, where faculty could lead them on walking tours rich in historical, architectural and cultural significance. Faculty also insisted, from time to time, on team dining in local restaurants. (See Figure 5).

Figure 4: Three robots remain standing (front, left) after the RoboCats are defeated by a team from Japan (front, right) in Padua. Co-author Cindy Marling is front and center, while co-author David Chelberg appears third from the left in the back row.

Figure 5: A Real Portuguese Restaurant

The second cultural problem was that RoboCup is inherently interdisciplinary, and the Mechanical Engineering (ME), Electrical Engineering (EE) and Computer Science (CS) team members were not used to working together. In Padua, the ME students complained that, since they didn't understand what the EE and CS students were doing, they didn't have much to contribute except when the robots broke down. We brainstormed together, MEs, EEs and CS folks, and proposed a solution to be implemented the following school year. We decided to offer a single course, *Intelligent Robotics via RoboCup*, and to cross-list it for all three disciplines, at both graduate and undergraduate levels. Besides formal lectures and group lab times, students conducted "mini-lessons" to share their disciplinary expertise, including, for example, an EE-led "How to Solder" exercise. The faculty members who team taught the course believed it to be a success, until the course evaluations were published. An ME student wrote, under cover of anonymity, "Unable to understand most of the EE/CS folk or contribute to discussions... consider moving MEs to separate meeting..."

Another culture clash occurred when we tried to address the problem of inadequate soccer expertise. The Good Old Fashioned AI (GOFAI) solution proposed was to enlist the aid of a human expert. We invited a talented soccer coach to meet with us and showed him our robots. We were unsure if he could afford a sufficient time commitment, but we were certain he would be excited by our grand challenge project. Instead, the horrified jock told the disconcerted nerds, "I just hope I never have to play soccer against robots."

Communications Failed

We were also plagued by communications failures. Robots and computers communicate via wireless radios. Our team's design used 802.11b wireless communications, as we had on-board processors running Linux, and our students were well-versed in this communications protocol. Many other teams use serial wireless radios operating at different frequencies, so we expected to be in a relatively unused part of the radio spectrum. This assumption proved to be *very* wrong. To understand just how wrong this assumption was, consider that at a RoboCup competition, there are typically thousands of students and faculty all with laptops, and various other gadgets that are capable of transmitting on the 2.4GHz frequency band used for the 802.11b communications protocol. Other leagues use 802.11b for communications as well, notably the Four-Legged League and RoboCup Rescue. In addition, although RoboCup's leadership tries to get people to turn off their laptops' wireless communications, many people either don't know how or forget to do this. The result is that there is a *lot* of interference at this frequency band. Furthermore, with the advent of Bluetooth technology, which shares the same frequency band, the interference has only increased.

In our first international competition in Japan, we had relatively few problems with communications, due primarily to the fact that the conference organizers hired engineers to install 802.11b base stations with directional antennas directly over each competition field. However, in Italy, we had significant trouble in getting our off-field computers communicating effectively with our robots. One suggested solution, found on the Web, was to use a Pringles potato crisps can to build an antenna that would potentially improve reception. We were able to find Pringles in a Padua grocery store, and our students were happy to eat the potato crisps, circumventing the ban on American food to obtain the empty cans. Unfortunately, however, we did not achieve success at improving communications with this approach. Some teams came prepared with "extreme" antennas that were optimized for their own communications systems, showing us that we were not alone in our communications difficulties. Our current robots are equipped with both 802.11a (a much less used frequency band), and a separate serial wireless system capable of using several different frequency bands.

Human-human communications failures occurred frequently when students from different engineering disciplines spoke the same words but meant different things. This is one reason we started our interdisciplinary course with ME, EE and CS students all working and talking together. Many graduating RoboCup students have told us that employers were impressed by the fact that they already had experience in dealing with these interdisciplinary issues in school.

Human-human communications failures occurred in non-technical contexts as well, when people from different cultures spoke the same words but meant different things. Some of the RoboCats team members asked the English-speaking concierge of our Japanese hotel to recommend a good place for them to meet and get to know some local girls. They were sent to the local red light district, much to their chagrin and the subsequent horror of their faculty advisors.

Autonomy had its Disadvantages

Asimov noted early on that autonomous robots could potentially be agents of evil as well as agents for good (Asimov 1950). The RoboCats robots were not exactly evil, but they did have minds of their own. As shown in Figure 6, one robot chose to dance when it should have been playing soccer. In Figure 7, another robot programmed a professor to play in the Four-Legged League.

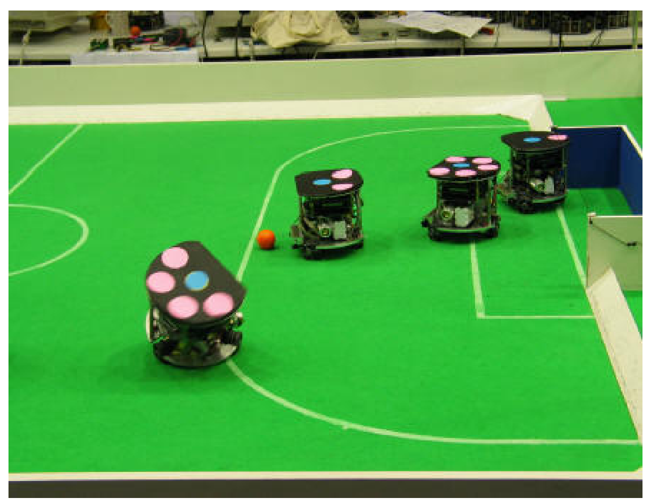

Figure 6: Robot Dances During Soccer Game

Why

Easy Explanations

The most obvious explanation is Murphy's Law. Things went wrong because they could. Another easy explanation is

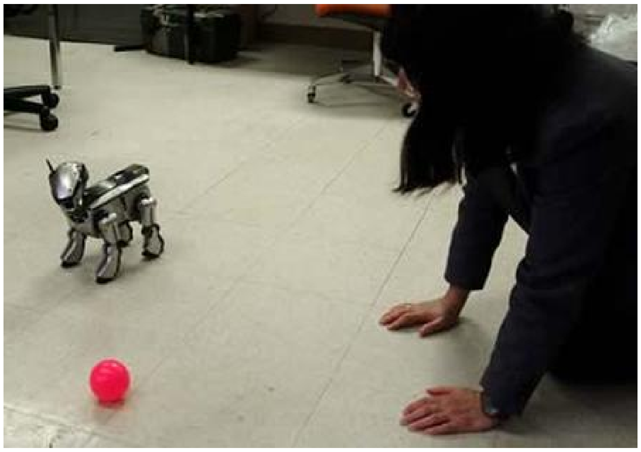

Figure 7: Robot Programs Professor for Four-Legged Play

that the considerable mechanical, athletic and cultural obstacles faced by the RoboCats were just too great to overcome. While these explanations hold some truth, they are not very satisfying.

Rank Speculation

We speculate that perfect progress is inversely correlated with research challenge. For example, Orville and Wilbur Wright conducted numerous unsuccessful experiments before they conquered flight. They built on the efforts of predecessors who literally crashed and died (State Library of North Carolina 2008). Failure might be viewed as a catalyst for progress. This view has not been well-received by the allocators of money for robot acquisitions and faculty salaries.

The Humans Couldn't Plan, Either

One of the toughest problems for AI planning systems is the frame problem (Shanahan 2006). The frame problem, first introduced in (McCarthy & Hayes 1969), involves reasoning about the things that change or stay the same when an action is taken. Many of the failures documented in this paper can be attributed to faulty assumptions about what changes due to a particular action. For example, we assumed that since CO_2 cartridges were readily available in the United States, they would also be available throughout the world. This assumption is akin to not understanding that the action of moving to an international venue to participate in RoboCup would have unanticipated consequences. Thus our human planning suffered from the frame problem. Classical AI approaches to dealing with the frame problem would not, in general, help in these types of situations, as humans often do not realize that they have even made an unwarranted assumption until a problem develops.

What can be done? One possible approach is to learn from what went wrong, to enable us to widen the scope of the anticipated consequences of actions. This approach suggests that we should build AI planning systems to learn from their mistakes, possibly through the use of CBR to classify and recount past unanticipated consequences of actions. Case-based planning (CBP) is an approach in which planning occurs through reusing past planning successes and avoiding past planning failures. As explained in (Hammond 1990), "new plans should be based on the planner's knowledge of what has succeeded and failed in the past." Unfortunately, we started out without any experience in moving teams of humans and robots to international venues or of winning robotic soccer matches once we got there. Our initially empty case base, which grew by only one case per year, precluded our effective use of CBP for RoboCup.

Another problem we had is that, occasionally, we had enough knowledge to predict possible failures, but we lacked sufficient resources in terms of people, time and money, to be able to avoid them. Once again, there is an AI planning analogue. In planning, resource constraints often preclude finding an optimal plan. We, as humans, find planning under resourse constraints extremely difficult, so it is not surprising that this is a challenge for AI research. With the advent of more and more autonomous machines, from RoboCup to unmanned vehicles, this becomes an even more important topic for AI research and practical AI applications.

On Knowing Why

One of the tenets of CBR is that past experience can improve future performance *even if you don't know why*. The classic example is the CLAVIER system, which was built to help Lockheed configure layouts for loading an autoclave, a large pressurized convection oven (Hennessy & Hinkle 1992). Lockheed engineers did not fully understand the physical properties of the autoclave, which, when improperly loaded, could ruin the parts it was intended to cure. They never did formalize the autoclave's curing process, but by remembering those layouts that produced properly cured parts, they improved the oven's yield. If knowing exactly *why* things go wrong remains a challenge, storing and recalling cases of What Went Wrong may still improve AI research results.

Epilogue

While Cindy Marling now focuses her research elsewhere (Bichindaritz & Marling 2006), David Chelberg remains team leader of the RoboCats. The RoboCats last competed at the American Open in Atlanta, Georgia, in 2005, where they finished second in the Small-Size League. The RoboCats are currently in a rebuilding phase as they raise funds for next year's competitions. Our latest robot features a new four-wheel design, an electrically powered solenoid kicker, and a totally redesigned controller board with on-board inertial guidance sensors. We expect to have many more stories of What Went Wrong and Why to relate in the future. We note, in conclusion, that failure is a subjective concept, open to human and/or machine interpretation. As Thomas Edison once said, "I have not failed. I've just found 10,000 ways that won't work."

Acknowledgements

We would like to thank the many Ohio University undergraduate students, graduate students, and faculty members who have embraced the RoboCup challenge over the past several years. It is with the deepest gratitude and not a little sadness that we finally acknowledge the generous nonagenarian Mrs. Beth K. Stocker, who financed the RoboCats' international travel and refused to be publicly acknowledged during her lifetime.

References

Asimov, I. 1950. *I, Robot*. Garden City, NY: Doubleday.

Bichindaritz, I., and Marling, C. 2006. Case-based reasoning in the health sciences: What's next? *Artificial Intelligence in Medicine* 36:127–135.

Gillen, M.; Lakshmikumar, A.; Chelberg, D.; Marling, C.; and Welch, L. 2002. A hybrid hierarchical schema-based architecture for distributed autonomous agents. In *Intelligent Distributed and Embedded Systems: Papers from the 2002 AAAI Spring Symposium*. Menlo Park, CA: AAAI Press.

Hammond, K. J. 1989. *Case-Based Planning: Viewing Planning as a Memory Task*. Boston, MA: Academic Press.

Hammond, K. J. 1990. Case-based planning: A framework for planning from experience. *Cognitive Science* 14:385–443.

Hennessy, D., and Hinkle, D. 1992. Applying case-based reasoning to autoclave loading. *IEEE Expert* 7(5):21–26.

Kolodner, J. 1993. *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufman.

Law, V. 2003. CMU robots roll to victory. *Pittsburgh Tribune-Review* pages B1 and B6. May 5, 2003, http://www.pittsburghlive.com/x/pittsburghtrib/s_132780.html, accessed February, 2008.

Marling, C.; Tomko, M.; Gillen, M.; Alexander, D.; and Chelberg, D. 2003. Case-based reasoning for planning and world modeling in the RoboCup small size league. In *Proceedings of the IJCAI-03 Workshop on Issues in Designing Physical Agents for Dynamic Real-Time Environments: World Modeling, Planning, Learning and Communicating*.

McCarthy, J., and Hayes, P. 1969. Some philosophical problems from the standpoint of artificial intelligence. In Michie, D., and Meltzer, B., eds., *Machine Intelligence 4*, 463–502. Edinburgh: Edinburgh University Press.

Shanahan, M. 2006. The frame problem. In Zalta, E. N., ed., *Stanford Encyclopedia of Philosophy (Spring 2006 Edition)*. http://plato.stanford.edu/archives/spr2006/entries/frame-problem/, accessed May, 2008.

State Library of North Carolina. 2008. North Carolina Encyclopedia, Wright brothers national memorial. http://statelibrary.dcr.state.nc.us/nc/ncsites/wright1.htm, accessed February, 2008.

Visser, U., and Burkhard, H. 2007. RoboCup: 10 years of achievements and future challenges. *AI Magazine* 28(2):115–132.